

**AFRL-IF-RS-TR-2003-21**  
**Final Technical Report**  
**February 2003**



# **CONTROL COORDINATION OF MULTIPLE AGENTS THROUGH DECISION THEORETIC AND ECONOMIC METHODS**

**Stanford University**

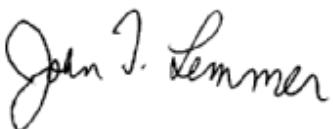
*APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.*

**AIR FORCE RESEARCH LABORATORY  
INFORMATION DIRECTORATE  
ROME RESEARCH SITE  
ROME, NEW YORK**

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2003-21 has been reviewed and is approved for publication.

APPROVED:



JOHN F. LEMMER  
Project Engineer

FOR THE DIRECTOR:



JAMES W. CUSACK, Chief  
Information Systems Division  
Information Directorate

# REPORT DOCUMENTATION PAGE

*Form Approved  
OMB No. 074-0188*

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503

<b>1. AGENCY USE ONLY (Leave blank)</b>			<b>2. REPORT DATE</b> FEBRUARY 2003		<b>3. REPORT TYPE AND DATES COVERED</b> Final Jun 98 – Jun 00	
<b>4. TITLE AND SUBTITLE</b> CONTROL COORDINATION OF MULTIPLE AGENTS THROUGH DECISION THEORETIC AND ECONOMIC METHODS			<b>5. FUNDING NUMBERS</b> C - F30602-98-C-0214 PE - 63760E PR - AGEN TA - T0 WU - 01			
<b>6. AUTHOR(S)</b> Yoav Shoham			<b>8. PERFORMING ORGANIZATION REPORT NUMBER</b>			
<b>7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)</b> Stanford University Office of Sponsored Research 320 Panama Street Stanford California 94305-4631						
<b>9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)</b> Air Force Research Laboratory/IFSF 525 Brooks Road Rome New York 13441-4505			<b>10. SPONSORING / MONITORING AGENCY REPORT NUMBER</b> AFRL-IF-RS-TR-2003-21			
<b>11. SUPPLEMENTARY NOTES</b>  AFRL Project Engineer: John F. Lemmer/IFSF/(315) 330-3657/ John.Lemmer@rl.af.mil						
<b>12a. DISTRIBUTION / AVAILABILITY STATEMENT</b> APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.					<b>12b. DISTRIBUTION CODE</b>	
<b>13. ABSTRACT (Maximum 200 Words)</b> The control and coordination of multi-agent systems is a major scientific and technological challenge. In settings where agents must act in flexible, hostile, distributed environments—such as those faced in military domains—the design of effective techniques for control coordination, competition, and adaptation becomes a task of great importance. There has been growing interest in the application of methods and approaches from economics to these problems; however, traditional economic methods lack many ingredients that are essential to make them useful for large-scale computational multi-agent systems. In our work we tackle some of these basic issues. In particular, we address: 1. the allocation of complementary and substitutable tasks to self-interested agents; 2. adaptation in hostile environments; 3. coordination for the assignment of a task among self interested bidders; 4. computationally-motivated representations of economic interactions; and 5. updating agents' beliefs after receiving new information. Our objective is therefore to introduce economic methods into the context of control and coordination of multi-agent systems, while generalizing and extending these methods to become efficient and effective. An important part of our approach is the identification and management of deep computational problems. We also present new theories which are essential for any flexible and dynamic practical multi-agent system.						
<b>14. SUBJECT TERMS</b> Combinatorial Auctions, Reinforcement Learning, Mechanism Design, Belief Revision					<b>15. NUMBER OF PAGES</b> 187	
					<b>16. PRICE CODE</b>	
<b>17. SECURITY CLASSIFICATION OF REPORT</b>  UNCLASSIFIED		<b>18. SECURITY CLASSIFICATION OF THIS PAGE</b>  UNCLASSIFIED		<b>19. SECURITY CLASSIFICATION OF ABSTRACT</b>  UNCLASSIFIED		
				<b>20. LIMITATION OF ABSTRACT</b>  UL		

## TABLE OF CONTENTS

Overview.....	1
1. Combinatorial Auctions.....	1
2. Adaptation in Multi-Agent Settings.....	2
3. Mechanism Design.....	2
4. Representation.....	4
5. Belief Revision and Belief Fusion .....	4
APPENDIX A Taming the Computational Complexity of Combinatorial Auctions: Optimal and Approximate Approaches .....	6
APPENDIX B An Algorithm for Multi-Unit Combinatorial Auctions .....	12
APPENDIX C Towards a Universal Test Suite for combinatorial Auction Algorithms.....	18
APPENDIX D R-Max – A General Polynominal Time Algorithm for Near-Optimal Reinforcement Learning .....	29
APPENDIX E Bidding Clubs: Institutionalized Collusion in Auctions.....	46
APPENDIX F Bidding Clubs in First-Price Auctions.....	53
APPENDIX G Mechanism Design with Incomplete Languages .....	80
APPENDIX H Anonymous Bidding and Revenue Maximization .....	116
APPENDIX I From Belief Revision to Belief Fusion .....	121
APPENDIX J Representing and Aggregating Conflicting Beliefs.....	156
APPENDIX K A Normative Examination of Ensemble Learning Algorithms.....	168
APPENDIX L Aggregating Learned Probabilistic Beliefs.....	176

# **Overview**

The control and coordination of multi-agent systems is a major scientific and technological challenge. When facing large-scale multi-agent settings where the agents are to act in flexible, hostile and distributed environments-such as those faced in military domains-the design of effective techniques for dealing with control, coordination, competition, and adaptation becomes a task of great importance. In recent years there has been growing interest in the application of methods and approaches from economics, for example the application of classic solutions from the theory of economic mechanism design to task allocation in non-cooperative dynamic environments. However, traditional economic methods lack many ingredients that are essential to make them applicable to large-scale computational multi-agent systems. In our work we tackle some of these basic issues. In particular, we address the allocation of complementary and substitutable tasks to self-interested agents, adaptation in hostile environments, coordination for the assignment of a task among self-interested bidders, computationally-motivated representations of economic interactions, and the updating of agents' beliefs after receiving new information. Our objective is therefore to introduce economic methods into the context of control and coordination of multi-agent systems, while generalizing and extending these methods to become efficient and effective. An important part of our approach is the identification and management of the deep computational problems which frequently arise in the control and coordination of large-scale multi-agent systems. We also present new theories which are essential for any flexible and dynamic practical multi-agent system.

Our work in the COABS project may be seen as addressing five classes of problems:

## **1. Combinatorial Auctions**

One primary economic mechanism upon which we chose to focus is the combinatorial auction. Combinatorial auctions involve the sale of multiple goods in a single auction, in cases where bidders' valuations may exhibit both complementarities (i.e., a bidder's willingness to pay for a bundle may exceed the sum of that bidder's valuation for each individual item in the bundle) and substitutability's (e.g., a bidder may be willing to win only one of a set of bundles). To allow bidders to express complementarities in their valuations, combinatorial auctions allow bidders to request "all-or-nothing" bundles of goods; bidders may also bid on subsets of these bundles if they are interested. To allow bidders to express substitutability's in their valuations, combinatorial auctions allow bidders to designate a set of bids as mutually exclusive-i.e., to indicate that only one of these bids is allowed to win, even if the seller would otherwise prefer to select more than one of these bids. Combinatorial auctions can lead to increased social welfare and/or seller revenue, but they come at a computational cost. Determining the set of winning bids in a combinatorial auction is an NP-hard computational problem. Nevertheless, we developed techniques to solve problems of interesting size by using a variety of different optimization techniques; we also investigated the design of test data for benchmarking such optimization algorithms. Our other research on combinatorial auctions included investigating bidder strategies when goods are allocated through sequential, single-good auctions, and an alternative mechanism that maintains incentive compatibility even though goods are not always allocated to the bidders willing to pay the most for them.

All of these papers, and the papers in the following sections, may be found in the appendix:

- *Sequential Auctions for the Allocation of Resources with Complementarities* (C. Boutilier, M. Goldszmidt and B. Sabata): presented at IJCAI-99.
- *Taming the Computational Complexity of Combinatorial Auctions: Optimal and Approximate Approaches* (Fujishima, Leyton-Brown, Shoham): presented at IJCAI- 99.
- *Incentive Compatibility in Rapid, Approximately Efficient Combinatorial Auctions* (D Lehmann, L. O'Callahan, and Y. Shoham): presented at the First ACM Conference on Electronic Commerce (EC'99)
- *An Algorithm for Multi-Unit Combinatorial Auctions* (K. Leyton-Brown, M. Tennenholtz, Y. Shoham): presented at Games-2000, AAAI-2000 and the International Symposium for Mathematical Programming (ISMP-2000).
- *Towards a Universal Test Suite for Combinatorial Auctions* (Leyton-Brown, Pearson, Shoham): EC-OO.

## 2. Adaptation in Multi-Agent Settings

Studying adaptation in multi agent settings was an important component of our research agenda. Indeed, the simultaneous adaptation of multiple agents has profound impact on the design of robust command and control methods. The phenomenon of adaptation in multi-agent systems is considerably different from adaptation in the single-agent case. This is true because the fact that multiple agents simultaneously adapt to each other implies that even simple adaptation rules can lead to complex behaviors. In order to tackle this issue we addressed the problem of reinforcement learning in various classes of stochastic games. Stochastic games extend upon and incorporate features of repeated games and Markov Decision Processes (MDPs), and are a very general model of multi-agent interaction. Our work on this topic had two main threads. First, we studied algorithms that could learn bidding policies in complex auction settings, and investigated the behavior of these algorithms. Second, we developed a reinforcement learning algorithm for stochastic games that finds near-optimal policies in polynomial time, and which also introduces a new approach for dealing with the exploration vs. exploitation tradeoff.

- *Continuous Value Function Approximation for Sequential Bidding Policies* (C. Boutilier, M. Goldszmidt and B. Sabata): presented at the Fifteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-99).
- *Conditional, Hierarchical Multi-Agent Preferences* (?): presented at the Seventh Conference on Theoretical Aspects of Rationality and Knowledge (TARK VII).
- *Sequential Optimality and Coordination in Multi-agent Systems* (C. Boutilier): presented at ?
- *R -max: A near optimal polynomial time reinforcement learning algorithm*, (Ronen Brafman and Moshe Tennenholtz): presented at UCAI'01.

## 3. Mechanism Design

One of the principal techniques for the control of multi-agent systems is the deployment

of an economic mechanism which will influence agents' behavior by giving them incentives for taking desirable actions. This *mechanism design* approach underlies a number of the research projects we undertook as part of our participation in the COABS project; because they are all so diverse we survey them individually here.

Ascending bid auctions--such as the familiar English-style auction of Sotheby and eBay--suffer the problem of being unpredictably long. This is unacceptable in mission critical, urgent applications, of the sort encountered in the military. The alternative--running a quick, one-shot sealed-bid auction--has the advantage of being fast, but unfortunately it does not possess the nice optimization properties of ascending-bid auctions in the presence of so-called common values. We were able to devise a novel auction mechanism, which combines the merits of both.

Finding ways of designing smart agents to assist bidders in auctions is fundamental to introducing agents' coordination to the context of economic mechanisms design. Our research emphasized protocols for coordinating groups of bidders through the paradigm of "bidding clubs"--groups of bidders who share information before participating in an auction, in such a way that all the members of a bidding club benefit. In our first paper on this topic we developed basic bidding club protocols for five fundamental auction settings; in our second paper we conducted a more rigorous and general theoretical analysis of bidding clubs in first-price auctions.

"Rational computation" presents a new model of computation based upon principles of rationality, which, we argue, are appropriate in a non-cooperative computing environment such as the Internet. In this work we developed a theory which looks at markets as computing devices and attempts to quantify their computing power.

Although VCG mechanisms have many appealing properties, their essential intractability prevents them from being used for complex problems like combinatorial auctions. We introduced a general way to overcome this intractability and proved its properties.

As we consider the use of auctions for resource allocation we must take into account the possibility--and in some cases virtual certainty--that agents will hide their true identities, so that it becomes impossible not only to know who is behind a given bid but even whether two different bids were submitted by the same bidder. This has profound effect on the outcome of the auction, as the bidders learn to manipulate the auction by using this anonymity feature. We were able to characterize the equilibria of some auctions in such settings, which provides the first step towards designing auctions that can withstand anonymity.

- *Speeding Up Ascending-Bid Auctions* (Y. Fujisima, D. McAdams, and Y. Shoham): presented at IJCAI-99.
- *Bidding Clubs: Institutionalized Collusion in Auctions* (K. Leyton-Brown, M. Tennenholz, and Y. Shoham): presented at Games-2000, EC-OO, Brown University.
- *Rational computation* (M. Tennenholz, and Y. Shoham): published in AII.
- *Bidding Clubs for First-Price Auctions* (Leyton-Brown, Shoham and Tennenholz): submitted to GEB.
- *Mechanism Design With Incomplete Languages* (Ronen): presented at EC-OI.
- *Anonymous bidding in auctions* (Yossi Feinberg and Moshe Tennenholz): submitted to GEB.

## 4. Representation

Bayesian networks-graphical representations of probability distributions that explicitly describe independences inherent in these distributions-revolutionized the field of probabilistic inference. By capturing the underlying structure of distributions, they allowed for algorithms that made inference tractable in practice. We have studied the possibility of finding structured representations for games which give similar tractability benefits. We began by studying possible ways of graphically representing utilities. The idea is that such representations capture structure inherent in the utility functions in the same way that Bayesian networks capture independences in probability distributions. Next, we introduced Game networks (G nets), a novel representation for multi-agent decision problems. Compared to other game-theoretic representations, such as strategic or extensive forms, G nets are more structured and more compact; more fundamentally, G nets constitute a computationally advantageous framework for strategic inference, as both probability and utility independencies are captured in the structure of the network and can be exploited in order to simplify the inference process. An important aspect of multi-agent reasoning is the identification of some or all of the strategic equilibria in a game; we presented original convergence methods for strategic equilibrium which can take advantage of strategic separabilities in the G net structure in order to simplify the computations. We introduced Multi-Agent Influence Diagrams (MAIDs), which generalize the familiar Bayesian Network generalization of (single-agent) influence diagrams to the multi-agent case. Finally, we developed a novel approach to computing all equilibria of a multi agent game, based on homotopy methods and closely related to simulated annealing used in AI.

- *Expected Utility Networks* (P. La Mura and Y. Shoham): presented at UAI'99.
- *Game Networks* (P. La Mura): presented at the Sixteenth Conference on Uncertainty in Artificial Intelligence (UAI'OO).
- *Probabilistic Models for Agents' Beliefs and Decisions* (B. Milch and D. Koller): I presented at UAI'QQ.
- *Simulated Annealing of Game Equilibria: A Simple Adaptive Procedure Leading to ~ Nash Equilibrium* (P. La Mura and M. Pearson): presented?

## 5. Belief Revision and Belief Fusion

Often we want to combine the expertise of multiple experts in hopes of coming up with information that improves on all their individual beliefs. We studied the problem of automating this process. We considered different common representations, both qualitative and quantitative, of sources' beliefs and studied how information about the sources' expertise can be used to combine their beliefs in rigorous, justified ways. Our initial focus in solving this problem was on the situation where agents' belief states are represented as qualitative binary relations over possible worlds. Such representations are common in the belief revision community to represent not only agents' beliefs, but their *counterfactual* beliefs as well, i.e., not only what they believe at the moment, but what they would believe if the situation were somewhat different.

We introduced a novel belief fusion operator that aggregates the beliefs of two agents, each informed by a subset of sources (strictly) ranked by reliability. In the process we

defined *pedigreed belief states*, which enrich standard belief states with the source of each piece of information. We noted that the fusion operator satisfies the invariants of idempotence, associativity, and commutativity. As a result, it can be iterated without difficulty. We also defined belief *diffusion*; whereas fusion generally produces a belief state with more information than is possessed by either of its two arguments, diffusion produces a state with less information.

We considered the problem of representing collective beliefs and aggregating these beliefs when there may be conflicting sources of equal rank. We described a way to construct the belief state of an agent informed by a set of sources of varying degrees of reliability, giving a simple set-theory-based operator for combining the information of multiple agents. We also described a computationally effective way of computing the resulting belief state.

Ensemble learning algorithms combine the results of several classifiers to yield an aggregate classification. We presented a normative evaluation of combination methods, applying and extending existing axiomatizations from Social Choice theory and Statistics. For the case of multiple classes, we showed that several seemingly innocuous and desirable properties are mutually satisfied only by a *dictatorship*. A weaker set of properties admit only the weighted average combination rule. We exemplified these theoretical results with experiments on stock market data, demonstrating how ensembles of classifiers can exhibit canonical voting paradoxes.

Finally, we shifted our attention to the problem of aggregating beliefs when they are represented as probabilistic distributions. We proposed a framework, in which we assumed that nature generates samples from a 'true' distribution and different experts form their beliefs based on the subsets of the data they have a chance to observe. We showed that the well-known aggregation operator LinOP is ideally suited for use in our framework, and proposed a LinOP-based learning algorithm, inspired by the techniques developed for Bayesian learning, which aggregates the experts' distributions represented as Bayesian networks.

- *From Belief Revision to Belief Fusion* (P. Maynard-Reid II and Y. Shoham): presented at the Third Conference on Logic and the Foundations of Game and Decision Theory (LOFT3).
- *Belief Fusion: Aggregating Pedigreed Belief States* (P. Maynard-Reid II and Y. Shoham): published in the Journal of Logic, Language, and Information.
- *Representing and Aggregating Conflicting Beliefs* (P. Maynard-Reid II and D. Lehmann): presented at the Seventh International Conference on Knowledge Representation and Reasoning (KR '00).
- *A Normative Examination of Ensemble Learning Algorithms* (D. Pennock, P. Maynard-Reid II, C. L. Giles, and E. Horvitz): presented at the Seventeenth International Conference on Machine Learning (ICML '00).
- *Aggregating Learned Probabilistic Beliefs* (P. Maynard-Reid II and U. Chajewska): presented at UAI'01.

## Taming the Computational Complexity of Combinatorial Auctions: Optimal and Approximate Approaches

**Yuzo Fujishima, Kevin Leyton-Brown and Yoav Shoham**

Computer Science Department, Stanford University, Stanford CA, 94305

*fujisima@ccs.mt.nec.co.jp* (visiting from NEC Corporation)

*kevinlb@cs.stanford.edu*

*shoham@cs.stanford.edu*

### Abstract

In combinatorial auctions, multiple goods are sold simultaneously and bidders may bid for arbitrary combinations of goods. Determining the outcome of such an auction is an optimization problem that is NP-complete in the general case. We propose two methods of overcoming this apparent intractability. The first method, which is guaranteed to be optimal, reduces running time by structuring the search space so that a modified depth-first search usually avoids even considering allocations that contain conflicting bids. Caching and pruning are also used to speed searching. Our second method is a heuristic, market-based approach. It sets up a virtual multi-round auction in which a virtual agent represents each original bid bundle and places bids, according to a fixed strategy, for each good in that bundle. We show through experiments on synthetic data that (a) our first method finds optimal allocations quickly and offers good anytime performance, and (b) in many cases our second method, despite lacking guarantees regarding optimality or running time, quickly reaches solutions that are nearly optimal.

### 1 Combinatorial Auctions

Auction theory has received increasing attention from computer scientists in recent years.<sup>1</sup> One reason is the explosion of internet-based auctions. The use of auctions in business-to-business trades is also increasing rapidly [Cortese and Stepanek, 1998]. Within AI there is growing interest in using auction mechanisms to solve distributed resource allocation problems. For example, auctions and other market mechanisms are used in network bandwidth allocation, distributed configuration design, factory scheduling, and operating system memory allocation

[Clearwater, 1996]. Market-oriented programming has been particularly influential [Wellman, 1993; Mullen and Wellman, 1996].

The value of a good to a potential buyer can depend on what other goods s/he wins. We say that there exists complementarity between goods  $g$  and  $h$  to bidder  $b$  if  $u_b(\{g,h\}) > u_b(\{g\}) + u_b(\{h\})$ , where  $u_b(G)$  is the utility to  $b$  of acquiring the set of goods  $G$ . If goods  $g$  and  $h$  were auctioned separately, it is likely that neither of the typically desired properties for auctions—efficiency and revenue maximization—would hold. One way to accommodate complementarity in auctions is to allow bids for combinations of goods as well as individual goods. Generally, auctions in which multiple goods are auctioned simultaneously and bidders place as many bids as they want for different bundles of goods are called combinatorial auctions<sup>2</sup>.

It is also common for bidders to desire a second good less if they have already won a first. We say that there exists substitutability between goods  $g$  and  $h$  to bidder  $b$  when  $u_b(\{g,h\}) < u_b(\{g\}) + u_b(\{h\})$ . A common example of substitutability is for a bidder to be indifferent between several goods but not to want more than one. In order to be useful, a combinatorial auction mechanism should provide some way for bidders to indicate that goods are substitutable.

Combinatorial auctions are applicable to many real-world situations. In an auction for the right to use railroad segments a bidder desires a bundle of segments that connect two particular points; at the same time, there may be alternate paths between these points and the bidder needs only one [Brewer and Plott, 1996]. Similarly, in the FCC spectrum auction bidders may desire licenses for multiple geographical regions at the same frequency band while being indifferent to which particular band they receive [Milgrom, 1998]. The same situation also occurs in military operations when multiple units each have several alternate plans and each plan may require a different bundle of resources.

---

<sup>1</sup> This material is based upon work supported by DARPA under the CoABS program, contract #F30602-98-C-0214, and by a Stanford Graduate Fellowship.

---

<sup>2</sup> Auctions in which combinatorial bidding is allowed are alternately called *combinatorial* and *combinational*.

While economics and game theory provide many insights into the potential use of such auctions, they have little to say about computational considerations. In this paper we address the computational complexity of combinatorial auctions.

## 2 The Complexity Problem

There has been much work in economics and game theory on designing combinatorial auctions. The Clarke-Groves-Vickrey mechanism (also known as the Generalized Vickrey Auction, or GVA) has been particularly influential [Mas-Colell et al., 1995; Varian, 1995]. It is beyond the scope of this paper to review such mechanisms, but they share a central problem: given a collection of bids on bundles, finding a set of non-conflicting bids that maximizes revenue. (A more precise definition is given in Section 3.) This problem is easily shown to be NP-complete<sup>3</sup> [Rothkopf et al., 1995].

Several methods have been conceived to cope with the computational complexity of combinatorial auctions, most aiming to ease the difficulty of finding optimal allocations. They can be classified into three categories based on the strategies they use.

One strategy is to restrict the degree of freedom of bidding to simplify the task of finding optimal allocations. Rothkopf et al. show that an optimal allocation can be found in polynomial time if (1) each bid contains no more than two goods; (2) for any two bids, either they are disjoint or one is a subset of the other; or (3) each bid contains only consecutive goods given a one-dimensional ordering of goods [Rothkopf et al., 1995].

Another strategy is to shift the burden of finding an optimal allocation to bidders. [Banks et al., 1989] and [Bykowsky et al., 1995] have reported a mechanism called AUSM in which non-winning bids are pooled in a stand-by queue. Bidders can combine their bids with other bids currently in the queue to form new allocations. A new allocation is adopted if it generates more revenue than the previously best allocation.

A third strategy is to attempt to find an optimal allocation but to be satisfied with a sub-optimal allocation when the expenditure of further resources becomes unacceptable. In other words, the optimality of the allocation is traded-off with the resources required, especially time.

In this paper we present two algorithms. The first is an anytime algorithm that attempts to exploit a problem's particular bid structure to reduce the size of the search. It also reduces search time by caching partial results and by pruning the search tree. The second algorithm uses a market-based approach to determine an acceptable allocation, although it is not guaranteed to find an optimal one. We then show results of experiments with synthetic data suggesting that these methods, though not provided with formal guarantees, appear to have surprisingly good per-

formance. Additionally, the market-based approach appears to produce allocations that are always optimal or nearly optimal.<sup>4</sup>

## 3 Precise Problem Statement

In this paper we propose two methods for finding desirable allocations based on bids submitted. We start by formally defining the optimization problem. Denote the set of goods by  $G$  and the set of non-negative real numbers by  $R^+$ . A bid  $b=(p_b, G_b)$  is an element of  $S = R^+ \times (2^G - \{\emptyset\})$ . Let  $B$  be a subset of  $S$ . A set  $F \subseteq B$  is said to be feasible if  $\forall b, c \in F, G_b \cap G_c = \emptyset$ . Denote the set of all feasible allocations for  $B$  by  $\Phi(B)$ . Further, let  $G(B) = \cup_{b \in B} G_b$  be the set of goods contained in the bids of  $B$ .

**[Problem]** Find an allocation  $W \in \Phi(B)$  such that  $\forall F \in \Phi(B), \sum_{b \in F} p_b \leq \sum_{b \in W} p_b$ . Such an allocation is said to be optimal or revenue maximizing.

What kind of value interrelation between goods can be represented by the bids defined above? Clearly, complementary values are easily accommodated. Suppose a bidder bids \$20 for each of  $\{g\}$  and  $\{h\}$ , and \$50 for  $\{g,h\}$ . In this case any revenue-maximizing algorithm will correctly select the  $\{g,h\}$  bid instead of  $\{g\}$  and  $\{h\}$ .

This bid format is also sufficient for representing substitutability through an encoding trick. Suppose a bidder is willing to pay \$20 for  $\{g\}$  and \$30 for  $\{h\}$  but only \$40 for  $\{g,h\}$ . In this case, bids cannot be submitted as before since the revenue-maximizing algorithm would select the pair  $\{g\}$  and  $\{h\}$  over  $\{g,h\}$ , charging the bidder \$50 instead of \$40 for  $g$  and  $h$ . However, this problem can be solved by the introduction of ‘dummy goods’—virtual goods that enforce an exclusive-or relationship. (Each dummy good must appear only in a single bidder’s bids.) In our example, the bidder could submit the following bids: (\$20,  $\{g,d\}$ ), (\$30,  $\{h,d\}$ ), and (\$40,  $\{g,h\}$ ) where  $d$  is a new, unique dummy good. The first two bids are now mutually exclusive and so will never be allocated together. This technique can lead to a combinatorial explosion in the number of bids if many goods are substitutable, but in many interesting cases this does not arise.

## 4 CASS Algorithm

When the number of goods and bids is small enough, an exhaustive search can be used to determine the optimal allocation. We propose an algorithm, Combinatorial Auction Structured Search (CASS), presented as a naïve brute-force approach followed by four improvements. CASS considers fewer partial allocations than the brute-force method because it structures the search space to avoid considering allocations containing conflicting bids. It also caches the results of partial searches and prunes the search tree. Finally, it may be used as an any-

<sup>3</sup> The GVA has the additional shortcoming of requiring bidders to submit an unreasonably large number of bids, but we do not address this issue here.

<sup>4</sup> We do not analyze the impact of the approximation on the equilibrium strategies in auction mechanisms such as GVA; we will address this issue in a future paper.

time algorithm, as it tends to find good allocations quickly.

#### 4.1 Brute-Force Algorithm

Suppose there are  $|G|$  goods 1, 2, ...,  $|G|$ , and  $|B|$  bids 1, 2, ...,  $|B|$ . First, bids that will never be part of an optimal allocation are removed. That is, if for bid  $b_k=(p_k, G_k)$  there exists a bid  $b_l=(p_l, G_l)$  such that  $p_l > p_k$  and  $G_l \subseteq G_k$ , then  $b_k$  is removed because it can always be replaced by  $b_l$ , increasing revenue. Then for each good  $g$ , if there is no bid  $b=(x, \{g\})$  a dummy bid  $b=(0, \{g\})$  is added.

Our brute-force algorithm examines all feasible allocations through a depth-first search. Let  $x$  be the first bid and  $y$  be the last bid. Our implementation follows:

1. If  $x$  does not conflict with the current allocation, add  $x$  to the current allocation
2. Increment  $x$
3. If more bids can be added to the allocation, go to 2.
4. Update best revenue and allocation observed so far.
5. If  $y$  is contained in the current allocation, remove it, set  $x=y+1$  and repeat from 2.
6. Decrement  $y$ .
7. If  $y$  is not the first bid, go to 5.

#### 4.2 Improvement #1: Bins

A great deal of unnecessary computation is avoided in the brute-force algorithm by checking whether bids conflict with the current allocation before they are added. However, work is still required to determine that a combination is infeasible and to move on to the next bid. It would be desirable to structure the search space to reduce the number of infeasible allocations that are considered in the first place.

We can reduce the number of infeasible allocations considered by sorting bids into bins,  $D_i$ , containing all bids  $b$  where good  $i \in G_b$  and for all  $j$  such that  $j \in [1, i-1], j \notin G_b$ . Rather than always trying to add each bid to our allocation, we add at most one bid from every bin since all bids in a given bin are mutually exclusive.

In fact, we can often skip bins entirely. While considering bin  $D_i$ , if we observe that good  $j > i$  is already part of the allocation then we do not need to consider any of the bids in  $D_j$ . In general, instead of considering each bin in turn, skip to  $D_k$  where  $k \notin G(F)$  and  $\forall i < k, i \in G(F)$ .

#### 4.3 Improvement #2: Caching

Let  $F_i$  be the partial allocation under consideration when  $D_i$  is reached during a search. Define  $C_i \subseteq G(F_i)$  where  $\forall j \in G(F_i), j > i \leftrightarrow j \in C_i$ . Note that there are many different partial allocations  $F_{i1}, F_{i2}$ , etc., that share the same  $C_i$ , and that if  $C_{i1} = C_{i2}$  then the search trees for  $F_{i1}$  and  $F_{i2}$  are identical beyond  $D_i$ . It is therefore possible to cache partial searches based on  $C_i$ . However, caching all possible values of  $C_i$  would require a cache of size  $2^{|G|-|i-1|}$ , which would quickly become infeasible. Therefore, we only cache when  $C_i$  includes no more than  $k$  goods, where  $k$  is a threshold defined at

runtime for each bin.  $D_i$  requires a cache of size  $\sum_{j=0}^k \binom{|G|-i}{j}$ .

#### 4.4 Improvement #3: Pruning #1

Performance can be improved by backtracking whenever a given search path is provably unable to lead to a new best allocation. We can prune whenever  $C(F_{i1}) \subset C(F_{i2})$  and  $p(F_{i2}) + p(\text{cache}(F_{i2})) \leq \text{bestAllocation}$ . In this case, the sum of the revenue from the cached path beyond  $F_{i1}$  and the revenue leading up to  $F_{i2}$  is less than the revenue from the best allocation seen so far. Since  $F_{i1}$  allocates a superset of the goods allocated in  $F_{i2}$  (thus overestimating revenue), a better allocation would not be found by expanding  $F_{i2}$ .

#### 4.5 Improvement #4: Pruning #2

We can also backtrack when it is provably impossible to add any bids to the current allocation to generate more revenue than the current best allocation. Before starting the search we calculate an overestimate of the revenue that can be achieved with each good,  $o(g) = \max_{b|g \in b} p(b)/|G_b|$ .  $o(g)$  is the largest average price per bid of bids containing good  $g$ . We backtrack at any point during the search with allocation  $F$  if  $p(F) + \sum_{g \notin F} o(g) \leq p(\text{best\_allocation})$ .

This technique is most effective when good allocations are found quickly. Finding good allocations quickly is also useful if a solution is required before the algorithm has completed (i.e., if CASS is used as an anytime algorithm). We have found that good allocations are found early in the search when the bids in each bin are ordered in descending order of average price per good. Similarly, the pruning technique is most effective when the unallocated goods are those with the lowest  $o(g)$  values. To achieve this, we reorder bins so that for any two bins  $i$  and  $j$ ,  $o(g_i) > o(g_j) \leftrightarrow i < j$ .

### 5 VSA Algorithm

Our second algorithm is called Virtual Simultaneous Auction (VSA). This market-oriented method was inspired by market-oriented programming [Wellman, 1993; Mullen and Wellman, 1996] and the simultaneous ascending auction [Milgrom, 1998]. VSA generates a virtual simultaneous auction from the bids submitted in a real combinatorial auction, then simulates the virtual auction to find a good allocation of goods in the real auction.

#### 5.1 Algorithm

First, a virtual simultaneous auction is generated based on the bids submitted in a real combinatorial auction. For each bid  $b=(p_b, G_b)$  a virtual bidder  $v_b$  is created. The virtual bidders compete in a virtual simultaneous auction that has multiple rounds. Each virtual bidder  $v_b$  tries to win all the goods in  $G_b$  for the price  $p_b$  on an all-or-nothing basis. The virtual auction starts with no goods allocated and the

prices of all goods set to zero. The simultaneous auction is repeated round by round until either an optimal allocation is found or a pre-set time deadline is reached. In the latter case the current best allocation is adopted as the final result.

Each round of VSA has three phases: the virtual auction phase, the refinement phase and the update phase. In the virtual auction phase each virtual bidder bids for the goods they want. Each individual good is allocated to the highest bidder. If a bidder succeeds in winning all desired goods, that bidder becomes a temporary winner. Otherwise the bidder becomes a temporary loser and returns all allocated goods to the auctioneer. In the refinement phase each of the losers is examined in a random order to see whether making that agent a temporary winner (and consequently making a different winner into a loser) would increase global revenue. If so, the list of winners is updated. Finally in the update phase the current highest price of each good is changed to reflect the price that its current winner bid. The current highest price for unallocated goods is reset to zero.

Virtual bidders in VSA follow a simple strategy. If a bidder was the temporary winner in the previous round, the bidder does not bid in the current round. Otherwise, agents calculate the sum of the current highest prices of the goods required. If the sum exceeds an agent's budget, the agent does not bid because the agent will not be able to acquire all the goods simultaneously. If the sum is less than the budget, the agent bids such that the surplus (budget - sum) is equally divided among the goods.

## 5.2 Properties

In certain circumstances, VSA will find an optimal allocation. Additionally, it is sometimes possible to detect if an optimal allocation has been found, allowing the virtual auction to end before the deadline.

**[Theorem]** If no virtual bidder bids in a round in the virtual auction, the current set of winners is optimal.

**[Proof]** Assume that no agents bid in a given round. Define the function that calculates the revenue of an allocation  $F$  by  $r(F)=\sum_{b \in F} p_b$  and let  $O$  denote the optimal set of winners. Split the current set of winners  $W$  into two parts  $O_1$  and  $W_2$  such that  $O_1=O \cap W$  and  $W_2=W \cap \neg O_1$ . Also split  $O$  into  $O_1$  and  $O_2$  such that  $O_1$  is defined as before and  $O_2=O \cap \neg O_1$ . Further, split  $G$  into  $G_1$  and  $G_2$  such that  $G_1=\cup_{b \in O_1} G_b$  and  $G_2=G \cap \neg G_1$ . By the assumption, for each currently losing bidder, the sum of the current highest prices of the goods needed exceeds the bidder's budget. This is especially true for bidders in  $O_2$ , i.e.,  $\forall b \in O_2 p_b < \sum_{g \in G_b} h_g$  where  $h_g$  is the current highest price of good  $g$ . It follows that  $r(O_2)=\sum_{b \in O_2} p_b \leq \sum_{b \in O_2} \sum_{g \in G_b} h_g \leq \sum_{g \in G_2} h_g = \sum_{b \in W_2} \sum_{g \in G_b} h_g = \sum_{b \in W_2} p_b = r(W_2)$ . (Remember that the minimum price of a good that is not allocated to any agent is zero and agents always bid their entire budgets.) The inequality means that  $W$  is optimal because  $r(O)=r(O_1)+r(O_2) \leq r(O_1)+r(W_2)=r(W)$ .

However, there is no guarantee that auctions will always

finish, even if an optimal allocation is found.

**[Theorem]** There exists a set of bids  $B$  such that at least one virtual bidder always bids in every round of the virtual auction no matter what bidding strategy is used.

**[Proof]** Suppose  $B=\{a,b,c\}$  where  $a=\{p_a, \{1,2\}\}$ ,  $b=\{p_b, \{2,3\}\}$ , and  $c=\{p_c, \{3,1\}\}$ . Suppose further that  $p_a < p_b + p_c$ ,  $p_b < p_c + p_a$ , and  $p_c < p_a + p_b$ . Because the real bids are mutually exclusive, at most one virtual bidder becomes the temporary winner. If none is winning,  $h_1=h_2=h_3=0$  and all the bidders bid in the current round. Assume here that bidder  $a$  is currently winning. Then  $h_1+h_2=p_a$  and  $h_3=0$ . Assume that neither  $b$  nor  $c$  bids in the current round. Then for each of  $b$  and  $c$ , the sum of the prices of goods needed must be larger than or equal to the budget, i.e.,  $h_2+h_3=h_2 \geq p_b$  and  $h_3+h_1=h_1 \geq p_c$ . This means that  $p_a = h_1+h_2 \geq p_b+p_c$  and contradicts  $p_a < p_b+p_c$ . This argument doesn't depend on the bidding strategy as long as an agent bids if and only if their budget exceeds the sum of the minimum prices of the goods needed.

It is this property that makes the refinement phase of VSA important. Consider the case  $B=B_1 \cup B_2 \cup \dots$  where  $\forall i,j G(B_i) \cap G(B_j)=\emptyset$ ,  $|B_i|=3$  and each  $B_i$  satisfies the condition from the proof above. If we omit the refinement phase then the winner in each subset changes every round except the case where there is no winner. Therefore, an optimal global allocation is examined only when in every subset the optimal winner is temporarily winning. Such synchronization is unlikely to occur unless the number of subsets is very small. The refinement phase causes the optimal winners to become the temporary winners in every round, leading to an optimal allocation even though it is not detected as optimal. (In some cases where  $\exists i,j G(B_i) \cap G(B_j) \neq \emptyset$  or  $|B_i| > 3$  an optimal allocation may be impossible to achieve regardless of the time limit.)

## 6 Experimental Evaluation

As we have not yet determined each algorithm's formal complexity characteristics we conducted empirical tests. We evaluated (1) how running time varies with the number of bids, and (2) how percentage optimality of the best allocation varies with time, given a particular bid distribution and a fixed number of bids and goods.

### 6.1 Assumptions and Parameters

The space of this problem is large. Roughly speaking it has three degrees of freedom: the number of goods, the number of bids and the distribution of bids. Most problematic among these is the distribution. Precisely because of the computational complexity of combinatorial auctions there is little or no real data available. In the absence of such data we tested our algorithms against bids drawn randomly from specific distributions.

Throughout the experiments we used the following two distribution functions to determine how often a bid for  $n$  goods appears. The first is binomial,  $f_b(n)=p^n(1-p)^{N-n}N!/(n!(N-n)!)$ ,  $p=0.2$ , in which the prob-

ability of each good being included in a given bid is independent of which other goods are included. The second distribution is of exponential form,  $f_e(n) = Ce^{-x/p}$ ,  $p=5$ , representing the case where a bid for  $n+1$  goods appears  $e^{-1/p}$  times less often than a bid for  $n$  goods. The prices of bids for  $n$  goods is uniformly distributed between  $[n(1-d), n(1+d)]$ ,  $d=0.5$ .

We do not present any experiments varying the number of goods in this paper because of space constraints. We found that for both CASS and VSA running time increased exponentially with the number of goods.

We ran our experiments on a 450MHz Pentium II with 256MB of RAM, running Windows NT 4.0. 30 MB of RAM was used for the CASS cache. All algorithms were implemented in C++.

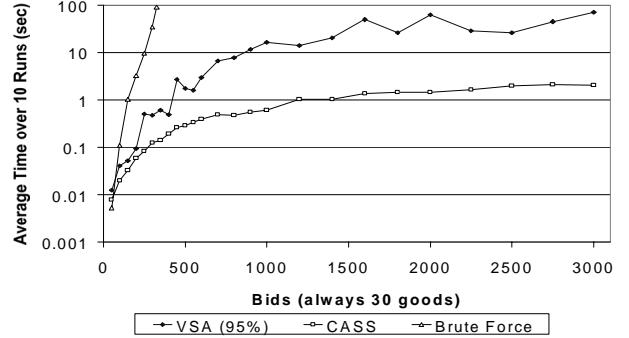
## 6.2 Results

To answer question (1) we measured the running time of CASS, VSA and the brute-force algorithm. Since VSA is not guaranteed to reach the optimal revenue, it was passed this value—calculated by CASS—and stopped when it found an allocation with revenue of at least 95% of optimal. All the results reported here are averages over 10 different runs. Figure 1 shows running time as a function of the number of bids with a binomial distribution, with the number of goods fixed at 30. Figure 2 shows the same thing for an exponential distribution, without the brute-force algorithm. To answer question (2), we measured the optimality of the output of both VSA and CASS as a function of time. Figure 3 shows both algorithms' performance with 15000 bids for 150 goods with a binomial distribution and Figure 4 shows 4500 bids for 45 goods with an exponential distribution.

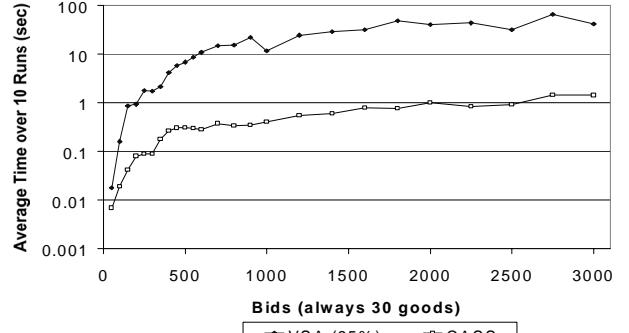
## 6.3 Discussion

CASS demonstrates excellent performance both in finding optimal allocations and as an anytime algorithm. In Figures 1 and 2 CASS remains roughly an order of magnitude faster than VSA as the number of bids increases. Both curves appear to grow sub-linearly on the logarithmic graph, suggesting polynomial-time performance. As the size of the problem is increased (Figures 3 and 4) CASS still performs better than VSA for the binomial distribution, but initially offers worse anytime performance for the exponential distribution. These results—and other experiments we have conducted—suggest that VSA is most likely to outperform CASS when the number of goods is relatively large compared to average bid length. (Note that VSA runs to a time limit, so the point at which VSA's curve ends is not meaningful.)

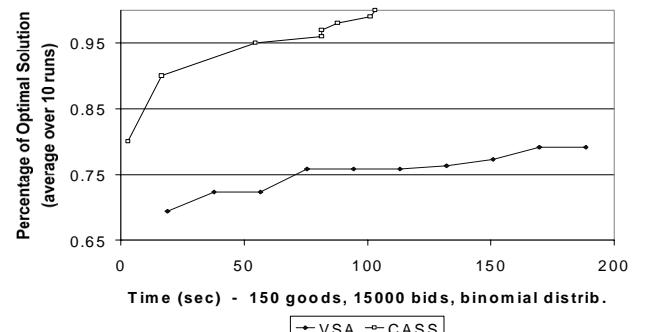
CASS's effectiveness is strongly influenced by the distribution of bids, particularly as the number of goods increases. If bids contain a large number of goods on average, improvement #1 will have a substantial effect because more bins will be skipped between every pair of bins that are considered, eliminating the need to individually examine all the bids in those bins. However, our



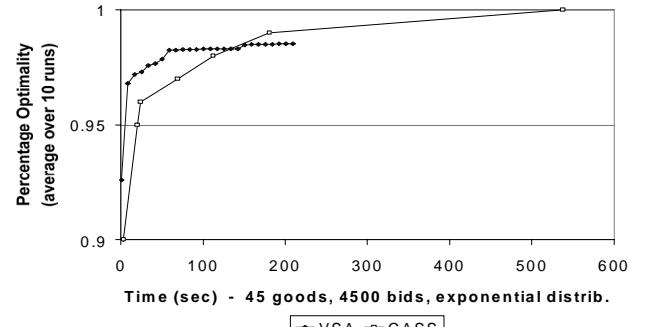
**Figure 1: Running Time Comparison (Binom. Dist.)**



**Figure 2: Running Time Comparison (Exp. Dist.)**



**Figure 3: Anytime Behavior (Binom. Dist.)**



**Figure 4: Anytime Behavior (Exp. Dist.)**

caching scheme favors distributions with small bids because they increase the likelihood that partial allocations will be cacheable. The pruning technique described in 4.4 reduces the number of nodes that are cached, lowering

memory consumption and making CASS feasible for larger problems. Our second pruning technique often improves performance by two orders of magnitude, though it is most effective when the variance of average price per bid is relatively small. This technique also reduces the optimal cache size, further reducing memory consumption. As a result of pruning, with pruning the amount of memory available for caching does not seem to be a limiting factor in CASS's performance.

VSA is interesting for two reasons. Firstly, it appears to offer good anytime performance in cases with small bids and many goods. Secondly, it provides a case study in the power of market-based optimization. Further work is needed to reach firm conclusions, but it appears that as a centralized optimization method VSA is overshadowed by other techniques. However, other attractions of market-based optimization—in particular its inherent distributed nature and robustness to change in problem specification—may make VSA attractive for some domains.

## 7 Related and Ongoing Work

As far as we are aware, the work most directly relevant to the ideas presented here is a paper by Sandholm [1999] that appears in these proceedings. Sandholm's Bidtree algorithm appears to be closely related to CASS, but important differences hold. In particular, Bidtree performs a secondary depth-first search to identify non-conflicting bids, whereas CASS's structured approach allows it to avoid considering most conflicting bids. Bidtree also performs no pruning analogous to our Improvement #3 and no caching. On the other hand, Bidtree uses an IDA\* search strategy rather than CASS's branch-and-bound approach, and does more preprocessing. We intend to continue studying the differences between these algorithms, including differences in experimental settings.

Our problem can of course be abstracted away from the auction motivation and viewed as a straightforward combinatorial optimization. This suggests a wealth of literature that could be applied. We are currently implementing some of these techniques and comparing them to our present results. We are especially interested in comparisons with mixed-integer programming and greedy methods. In particular, we have been investigating a new algorithm<sup>5</sup> that orders bids in descending order according to average price per good, and does a depth-first search with extensive pruning. This algorithm appears to offer performance similar to CASS, and we intend to report on it in a follow-up paper.

## 8 Conclusion

We have proposed two novel algorithms to mitigate the computational complexity of combinatorial auctions.

CASS determines optimal allocations very quickly, and also provides good anytime performance. In the future we

intend to pursue a formal analysis of CASS's computational complexity, and to test both CASS and VSA with data collected from real bidders.

VSA can determine near-optimal allocations even in cases with hundreds of goods and tens of thousands of bids. Since it has been infeasible to run CASS on much larger problems we do not yet know how close VSA comes to optimality in these cases. An investigation of VSA's limits remains an area for future work.

## References

- [Banks, et al., 1989] Jeffrey S. Banks, John O. Ledyard, and David P. Porter. Allocating uncertain and unresponsive resources: an experimental approach. *RAND Journal of Economics*, 20(1): 1-25, 1989.
- [Brewer and Plott, 1996] P.J. Brewer and C.R. Plott. A binary conflict ascending price (BICAP) mechanism for the decentralized allocation of the right to use railroad tracks. *International Journal of Industrial Organization*, 14:857-886, 1996]
- [Bykowsky et al., 1995] Mark M. Bykowsky, Robert J. Cull, and John O. Ledyard. Mutually destructive bidding: the FCC auction design problem. Social science working paper 916, California Institute of Technology, 1995.
- [Clearwater, 1996] Scott H. Clearwater, editor. *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, 1996.
- [Cortese and Stepanek, 1998] Amy E. Cortese and Marcia Stepanek. Good-bye to fixed pricing? *Business Week*, pages 71-84, May 4, 1998.
- [Mas-Colell et al., 1995] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.
- [Milgrom, 1998] Paul Milgrom. Putting auction theory to work: the simultaneous ascending auction. Working paper 98-002, Dept. Economics, Stanford University, 1998.
- [Mullen and Wellman, 1996] Tracey Mullen and Michael P. Wellman. Some issues in the design of the market-oriented agents. In M.J. Wooldridge, J.P. Muller, and M. Tambe, editors, *Intelligent Agents II: Agent Theories, Architectures, and Languages*. Springer-Verlag, 1996.
- [Rothkopf et al., 1995] Michael H. Rothkopf, Aleksandar Pekeč, and Ronald M. Harstad. Computationally manageable combinatorial auctions. DIMACS Technical Report 95-09, April 1995.
- [Sandholm, 1999] Tuomas Sandholm. An algorithm for optimal winner determination in combinatorial auctions. Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-99), Stockholm, 1999.
- [Varian, 1995] Hal R. Varian. Economic mechanism design for computerized agents. In *Proceedings of the First Usenix Conference on Electronic Commerce*, New York, July 1995.
- [Wellman, 1993] Michael P. Wellman. A market-oriented programming environment and its application to distributed multicommmodity flow problems. *Journal of Artificial Intelligence Research*, 1:1-23, 1993.

---

<sup>5</sup> This ongoing work is joined by Liadan O'Callaghan and Daniel Lehmann.

## APPENDIX B

# An Algorithm for Multi-Unit Combinatorial Auctions

**Kevin Leyton-Brown and Yoav Shoham and Moshe Tennenholtz**

Computer Science Department  
Stanford University, Stanford, CA 94305

### Abstract

We present a novel algorithm for computing the optimal winning bids in a combinatorial auction (CA), that is, an auction in which bidders bid for bundles of goods. All previously published algorithms are limited to single-unit CAs, already a hard computational problem. In contrast, here we address the more general problem in which each good may have multiple units, and each bid specifies an unrestricted number of units desired from each good. We prove the correctness of our branch-and-bound algorithm, which incorporates a specialized dynamic programming procedure. We then provide very encouraging initial experimental results from an implemented version of the algorithm.<sup>1</sup>

### Introduction

Auctions are the most widely studied mechanism in the mechanism design literature in economics and game theory (Fudenberg & Tirole 1991). This is due to the fact that auctions are basic protocols, serving as the building blocks of more elaborated mechanisms. Given the wide popularity of auctions on the Internet and the emergence of electronic commerce, where auctions serve as the most popular game-theoretic mechanism, efficient auction design has become a subject of considerable importance for researchers in multi-agent systems (e.g. (Wellman *et al.* 1998; Monderer & Tennenholtz 2000)). Of particular interest are multi-object auctions where the bids name bundles of goods, called combinatorial auctions (CA). For example, imagine an auction of used electronic equipment. A bidder may wish to bid  $x$  for a particular TV and  $y$  for a particular VCR, but  $z \neq x + y$  for the pair. In this example all the goods at auction are different, so we call the auction a single-unit CA. In contrast, consider an electronics manufacturer auctioning 100 identical TVs and 100 identical VCRs. A retailer who wants to buy 70 TVs and 30 VCRs would be indifferent between all bundles having 70 TVs and 30 VCRs. Rather than

having to bid on each of the  $\binom{100}{70} \cdot \binom{100}{30}$  distinct bundles, she would prefer to place the single bid (price,  $\{70 \text{ TVs}, 30 \text{ VCRs}\}$ ). We call an auction that allows such a bid a multi-unit CA.

In a combinatorial auction, a seller is faced with a set of price offers for various bundles of goods, and his aim is to allocate the goods in a way that maximizes his revenue. This optimization problem is intractable in the general case, even when each good has only a single unit (Rothkopf, Pekec, & Harstad 1998). Given this computational obstacle, two parallel lines of research have evolved. The first exposes tractable sub-cases of the combinatorial auctions problem. Most of this work has concentrated on identifying bidding restrictions that entail tractable optimization; see (Rothkopf, Pekec, & Harstad 1998; Nisan 1999; Tennenholtz 2000; Vries & Vohra 2000). Also, the case of infinitely divisible goods may be tractably solved by linear programming techniques. The other line of research addresses general combinatorial auctions. Although this is a class of intractable problems, in practice it is possible to address interestingly-large datasets with heuristic methods. It is desirable to do so because many economic situations are best modeled by a general CA, and bidders' strategic behavior is highly sensitive both to changes in the auction mechanism and to approximation of the optimal allocation (Nisan & Ronen 2000). Previous research on the optimization of general CA problems has focused exclusively on the simpler single-unit CA (Fujishima, Leyton-Brown, & Shoham 1999; Sandholm 1999; Lehmann, O'Callaghan, & Shoham 1999)). The general multi-unit problem has not previously been studied, nor have any heuristics for its solution been introduced.

In this paper we present a novel algorithm, termed CAMUS (Combinatorial Auction Multi-Unit Search), to compute the winners in a general, multi-unit combinatorial auction. A generalization and extension of our CASS algorithm for winner determination in single-unit CA's (Fujishima, Leyton-Brown, & Shoham 1999), CAMUS introduces a novel branch-and-bound technique that makes use of several additional procedures. A crucial component of any

Copyright © 2002, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup>This work was partly supported by DARPA grant number F30602-98-C-0214-P00005.

such technique is a function for computing upper bounds on the optimal outcome. We present such an upper bound function, tailored specifically to the multi-unit combinatorial auctions problem. We prove that this function gives an upper bound on the optimal revenue, which enables us to show that CAMUS is guaranteed to find optimal allocations. We also introduce dynamic programming techniques to more efficiently handle multi-unit single-good bids. In addition, we present techniques for pre-processing and caching, and heuristics for determining search orderings, further capitalizing on the inherent structure of multi-unit combinatorial auctions.

In the next section we formally define the general multi-unit combinatorial auction problem. In Section 3 we describe CAMUS. In Section 4 we deal in some more detail with some of CAMUS’s techniques. Due to lack of space, we cannot present all the CAMUS procedures in detail; however, this section will clarify its most fundamental components. In Section 5 we present our experimental setup and some experimental results.

## Problem Definition

We now define the computational problem associated with multi-unit combinatorial auctions.

Let  $G = \{g_1, g_2, \dots, g_m\}$  be a set of goods. Let  $q(j)$  denote the number of available units of good  $j$ . Consider a set of bids  $B = \{b_1, \dots, b_n\}$ . Bid  $b_i$  is a pair  $(p(b_i), e(b_i))$  where  $p(b_i)$  is the price offer of bid  $b_i$ , and  $e(b_i) = (e(b_i)_1, e(b_i)_2, \dots, e(b_i)_m)$  where  $e(b_i)_j$  is the number of requested units of good  $j$  in  $b_i$ . If there is no bid requesting  $k$  units of good  $i$  and 0 units of all goods  $j \neq i$  (for some  $1 \leq i \leq m$  and some  $1 \leq k \leq q(i)$ ) then, w.l.o.g, we augment  $B$  with a bid of price 0 for that bundle. An allocation  $\pi \subseteq B$  is a subset of the bids where  $\sum_{b \in \pi} e(b)_j \leq q(j)$  ( $1 \leq j \leq m$ ). A partial allocation  $\pi_{partial}$  is an allocation where, for some  $j$ ,  $\sum_{b \in \pi_{partial}} e(b)_j < q(j)$ . A full allocation is an allocation that is not partial. Let  $\Pi$  denote the set of all allocations. The multi-unit combinatorial auction problem is the computation of an optimal allocation, that is,  $\text{argmax}_{\pi \in \Pi} \sum_{b \in \pi} p(b)$ . In short, we are searching for a subset of the bids that will maximize the seller’s revenue while allocating each available unit at most once.

Note that the definition of the optimal allocation assumes that bids are additive—that an auction participant who submits multiple bids may be allocated any number of these bids for a price that equals the sum of each allocated bid’s price offer. In some cases, however, a participant may wish to submit two or more bids but require that at most one will be allocated. We permit such additional constraints through the use of *dummy goods*, introduced already in (Fujishima, Leyton-Brown, & Shoham 1999). Dummy goods are normal single-unit goods which do not correspond to actual goods in the auction, but serve to enforce mutual exclusion between bids. For example, if bids  $b_1$  and  $b_2$  referring to bundles

$e(b_1)$  and  $e(b_2)$  are intended to be mutually exclusive, we add a dummy good  $d$  to each bid:  $e(b_1)$  becomes  $e(b_1) \cup d$ , and  $e(b_2)$  becomes  $e(b_2) \cup d$ . Since the good  $d$  can be allocated only once, at most one of these bids will be in any allocation. (More generally, it is possible to introduce  $n$ -unit dummy goods to enforce the condition that no more than  $n$  of a set of bids may be allocated.) While dummy goods increase the expressive power of the bidding language, their use has no impact on the optimization algorithm. Hence, in the remainder of this paper we do not discriminate between dummy goods and real goods, and we assume that all bids are additive.

In the sequel, we will also make use of the following notation. Given an allocation  $\pi$  and a good  $i$ , we will denote the total number of units allocated in  $\pi$ , and the total number of units of good  $i$  allocated in  $\pi$ , by  $units(\pi)$  and  $units_i(\pi)$  respectively. In addition  $units(total)$  will denote the total number of units over all goods.

## Algorithm Definition

### Branch-and-Bound Search

Given a set of bids, CAMUS systematically compares the revenue from all full allocations in order to determine the optimal allocation. This comparison is implemented as a depth-first search: we build up a partial allocation one bid at a time. Once we have constructed a full allocation we backtrack, removing the most recently added bid from the partial allocation and adding a new bid instead. Sometimes we can safely *prune* the search tree, backtracking before a full allocation has been constructed. Every time a bid is added to the current allocation, CAMUS computes an estimate of the revenue that will be generated by the unallocated goods which remain. Provided that this estimate function  $o()$  always provides an upper bound on the actual revenue, we can prune whenever  $p(\pi) + o(\pi) \leq p(\pi_{best})$ , where  $\pi$  is the current allocation,  $p(\pi) = \sum_{b \in \pi} p(b)$  and  $\pi_{best}$  is the best allocation observed so far.

### Bins

*Bins* are partitioned sets of bids. Consider some ordering of the goods. There is one bin for each good, and each bid belongs to the bin corresponding to its lowest-order good. During the search we start in the first bin and consider adding each bid in turn. After adding a bid to our partial allocation we move to the bin corresponding to the lowest-order good with any unallocated units. For example, if the first bid we select requests all units of goods 1, 2 and 4, we next proceed to bin 3. Besides making it easy to avoid consideration of conflicting bids, bins are powerful because they allow the pruning function to consider context without significant computational cost. If bids in  $bin_i$  are currently being considered then the pruning function must only take into account bids from  $bin_i \dots bin_m$ . Because the partitioning

of bids into bins does not change during the search we may compute the pruning information for each bin in a preprocessing step.

## Subbins

In the multi-unit setting, we will often need to select more than one bid from a given bin. This leads to the idea of *subbins*. A subbin is a subset of the bids in a bin that is constructed during the search. Since subbins are created dynamically they cannot provide precomputed contextual information; rather, they facilitate the efficient selection of multiple bids from a given bin. Every time we add a bid to our partial allocation we create a new subbin containing the next set of bids to consider. If the search moves to a new bin, the new subbin is generated from the new bin by removing all bids that conflict with the current partial allocation. If the search remains in the same bin, the new subbin is created from the current subbin by removing conflicting bids as above, and additionally: if  $bid_1, bid_2, \dots, bid_i$  is the ordered set of elements in the current subbin and  $bid_j$  is the bid that was just chosen, then we remove all  $bid_k, k \leq j$ . In this way we consider all combinations of non-conflicting bids in each bin, rather than all permutations.

## Dominated Bids

Some bids may be removed from consideration in a polynomial-time preprocessing step. For each pair of bids  $(b_1, b_2)$  where both name the same goods but  $p(b_1) \geq p(b_2)$  and  $e(b_1)_j \leq e(b_2)_j$  for every good  $j$ , we may remove  $b_2$  from the list of bids to be considered during the search, as  $b_2$  is never preferable to  $b_1$  (hence we say that  $b_1$  *dominates*  $b_2$ ). However, it is possible that an optimal allocation contains both  $b_1$  and  $b_2$ . For this reason we store  $b_2$  in a secondary data structure associated with  $b_1$ , and consider adding it to an allocation only after adding  $b_1$ .

## Dynamic Programming

Singleton bids (that is, bids that name units from only one good) deserve special attention. These bids will generally be among the most computationally expensive to consider—the number of nodes to search after adding a very short bid is nearly the same as the number of nodes to search after skipping the bid, because a short bid allocates few units and hence conflicts with few other bids. Unfortunately, we expect that singleton bids will be quite common in a variety of real-world multi-unit CA's. CAMUS simplifies the problem of singleton bids by applying a polynomial-time dynamic programming technique as a preprocessing step. We construct a vector  $singleton_g$  for each good  $g$ , where each element of the vector is a set of singleton bids naming only good  $g$ .  $singleton_g(j)$  evaluates to the revenue-maximizing set of singleton bids totaling  $j$  units of good  $g$ . This frees us from having to consider singleton bids individually; instead, we consider only elements of the single-

ton vector and treat these elements as atomic bids during the search. Also, there is never a need to add more than one element from each singleton vector. To see why, imagine that we add both  $singleton_g(j)$  and  $singleton_g(k)$  to our partial allocation. These two elements may have bids in common, and additionally there may be singleton bids with more than  $\max(j, k)$  elements that would not conflict with our partial allocation but that we have not considered. Clearly, we would be better off adding the single element  $singleton_g(j + k)$ .

## Caching

Consider a partial allocation  $\pi_1$  that is reached during the search phase. If the search proceeds beyond  $\pi_1$  then  $o(\pi_1)$  was not sufficiently small to allow us to backtrack. Later in the search we may reach an allocation  $\pi_2$  which, by combining different bids, covers exactly the same number of units of the same goods as  $\pi_1$ . CAMUS incorporates a mechanism for caching the results of the search beyond  $\pi_1$  to generate a better estimate for the revenue given  $\pi_2$  than is given by  $o(\pi_2)$ . (Since  $\pi_1$  and  $\pi_2$  do not differ in the units of goods that remain,  $o(\pi_1) = o(\pi_2)$ .) Consider all the allocations extending  $\pi_1$  upon consideration of which the algorithm backtracked, denoted  $s_1, s_2, \dots, s_f$ . When we backtracked at each  $s_i$  we did so because  $p(s_i) + o(s_i) \leq p(\pi_{best})$ , as explained above. It follows that  $\max_i(p(s_i) + o(s_i))$  is an overestimate of the revenue attainable beyond  $\pi_1$ , and that it is a smaller overestimate than  $o(\pi_1)$  (if it were not, we would have backtracked at  $\pi_1$  instead). Since in general  $p(\pi_1) \neq p(\pi_2)$ , we cache the value  $\max_i(p(s_i) + o(s_i)) - p(\pi_1)$  and backtrack when  $p(\pi_2) + \text{cache}(\pi_2) \leq p(\pi_{best})$ . Our cache is implemented as a hash table, since caching is only beneficial to the overall search if lookup time is inconsequential. A consequence of this choice of data structure is that cache data may sometimes be overwritten; we overwrite an old entry in the cache when the search associated with the new entry examined more nodes. Even when we do overwrite useful data the error is not catastrophic, however: in the worst case we must simply search a subtree that we might otherwise have pruned.

## Heuristics

Two ordering heuristics are used to improve CAMUS's performance. First, we must determine an ordering of the goods; that is, which good corresponds to the first bin, which corresponds to the second, etc. For each good  $i$  we compute  $score_i = \frac{\text{numbids}_i \cdot q(i)}{\text{avgunits}_i}$ , where  $\text{numbids}_i$  is the number of bids that request good  $i$  and  $\text{avgunits}_i$  is the average number of *total* units (i.e., not just units of good  $i$ ) requested by these bids. We designate the lowest-order good as the good with the lowest score, then we recalculate the score for the remaining goods and repeat. The intuition behind this heuristic is as follows:

- We want to minimize the number of bids in low-order bins, to minimize early branching and thus to make each individual prune more effective.
- We want to minimize the number of units of goods corresponding to low-order bins, so that we will move quickly beyond the first few bins. As a result, the pruning function will be able to take into account more contextual information.
- We want to maximize the total number of units requested by bids in low-order bins. Taking these bids moves us more quickly towards the leaves of the search tree, again providing the pruning function with more contextual information.

Our second heuristic determines the ordering of bids within bins. Given current partial allocation  $\pi$ , we sort bids in a given bin in descending order of  $score(b_j)$ , where  $score(b_j) = \frac{p(b_j)}{units(b_j)} + o(\pi \cup b_j)$ . The intuition behind this heuristic is that the average price per unit of  $b_j$  is a measure of how promising the bid is, while the pruning overestimate for  $o(\pi \cup b_j)$  is an estimate of how promising the unallocated units are, given the partial allocation. This heuristic helps CAMUS to find good allocations quickly, improving anytime performance and also increasing  $\pi_{best}$ , making pruning more effective. Because the pruning overestimate depends on  $\pi$ , this ordering is performed dynamically rather than as a pre-processing step.

## CAMUS Outline

Based on the above, it is now possible to give an outline of the CAMUS algorithm:

- Process dominated bids.
- Determine an ordering on the goods, according to the good-ordering heuristic.
- Using the dynamic programming technique, determine the optimal combination of singleton bids totaling  $1 \dots q(j)$  for each good  $j$ .
- Partition all non-singleton bids into bins, according to the good ordering.
- Precompute pruning information for each bin.
- Set  $i = 1$  and  $\pi = \{\}$ .
- Recursive entry point:
  - For  $j = 1 \dots \text{number of bids in the current subbin of } bin_i$ .
    - \*  $\pi = \pi \cup bid_j$ .
    - \* If  $(p(\pi) + cache(\pi) \leq p(\pi_{best}))$  backtrack.
    - \* If  $(p(\pi) + o(\pi) \leq p(\pi_{best}))$  backtrack.
    - \* If  $(units(\pi) = units(total))$  record  $\pi$  if it is the best; backtrack.

- \* Set  $i$  to the index of the lowest-order good in  $\pi$  where  $units_i(\pi) < q(i)$ . ( $i$  may or may not change)
- \* Construct a new subbin based on the previous subbin of  $bin_i$  (which is  $bin_i$  itself if  $i$  changed above):
  - Include all  $bid_k$  from current subbin, where  $k > j$ .
  - Include all dominated bids associated with  $bid_j$ .
  - Include  $singleton_i(q(i) - units_i(\pi))$ .
  - Sort the subbin according to the subbin-ordering heuristic.
  - Recurse to the recursive entry point, above, and search this new subbin.
- \*  $\pi = \pi - bid_j$ .
- End For
- Return the optimal allocation:  $\pi_{best}$ .

## CAMUS procedures: a closer look

In this section we examine two of CAMUS's fundamental procedures more formally. Additional details will be presented in our full paper.

### Pruning

In this subsection we explain the implementation of CAMUS's pruning function and demonstrate that it is guaranteed not to underestimate the revenue attainable given a partial allocation. Consider a point in the search where we have constructed some partial allocation  $\pi$ . The task of our pruning function is to give an upper bound on the optimal revenue attainable from the unallocated items, using the remaining bids (i.e., the bids that may be encountered during the remainder of the search). Hence, in the sequel when we refer to goods, the number of units of a good and bids, we refer to what remains at our point in the search.

First, we provide an intuitive overview. For every (remaining) good  $j$  we will calculate a value  $v(j)$ . Simplifying slightly, this value is the largest average price per unit of all the (remaining) bids requesting units of good  $j$  that do not conflict with  $\pi$ , multiplied by the number of (remaining) units of  $j$ . The sum of  $v(j)$  values for all goods is an upper bound on optimal revenue because it relaxes the constraint that the bids in the optimal allocation may not conflict.

More formally, let  $G = \{g_1, g_2, \dots, g_m\}$  be a set of goods. Let  $q'(j)$  denote the number of available units of good  $j$ . Consider a set of bids  $B = \{b_1, \dots, b_n\}$ . Bid  $b_i$  is associated with a pair  $(p(b_i), e(b_i))$  where  $p(b_i)$  is the price offer of bid  $b_i$ , and  $e(b_i) = (e(b_i)_1, e(b_i)_2, \dots, e(b_i)_m)$  where  $e(b_i)_j$  is the requested number of units of good  $j$  in  $b_i$ . For each bid  $b_i$ , let  $a(b_i) = \frac{p(b_i)}{\sum_{1 \leq j \leq m} e(b_i)_j}$  be the average price per unit of bid  $b_i$ . Notice that the average price per unit may change dramatically from bid to bid, and it is a non-trivial notion; our technique will work for any arbitrary average price per unit. Let  $L(j)$  be a sorted list of the bids that refer to non-zero units of good  $j$ ; the list is sorted in a monotonically decreasing manner according to the  $a_i$ 's. Let  $|L(j)|$  denote the number of elements in  $L(j)$ , and let  $L(j)_k$  denote the  $k$ -th element of  $L(j)$ .

$v(j)$  is determined by the following algorithm:

```

Let  $v(j) := 0$ ;
Let  $m(j) := 0$ ;
For  $i := 1$  to  $|L(j)|$  do
if  $m(j) < q'(j)$  then
{let  $d := \min(e(L(j)_i), q(j) - m(j))$ ;  $m(j) = m(j) + d$ ;  $v(j) = v(j) + a(L(j)_i) \cdot d$ }

```

**Theorem 1** Let  $B^o = \{b_1^0, b_2^0, \dots, b_s^0\}$  be the bids in an optimal allocation. Then,  $R^o = \sum_{b \in B^o} p(b) \leq \sum_{1 \leq j \leq m} v(j)$ .

**Sketch of proof:** Consider the bid  $b^o \in B^o$ . Then,  $p(b^o) = \sum_{1 \leq j \leq m} a(b^o) \cdot e(b^o)_j$ . Hence,  $R^o = \sum_{b \in B^o} p(b) = \sum_{b \in B^o} \sum_{1 \leq j \leq m} a(b) \cdot e(b)_j$ . By changing the order of summation we get that  $R^o = \sum_{1 \leq j \leq m} \sum_{b \in B^o} a(b) \cdot e(b)_j$ . Notice that, given a particular  $j$ , the contribution of bid  $b$  to  $\sum_{b \in B^o} a(b) \cdot e(b)_j$  is  $a(b) \cdot e(b)_j$ . Recall now that  $v(j)$  has been constructed from the set of all bids that refer to good  $j$  by choosing the maximal available units of good  $j$  from the bids in  $L(j)$ , where these bids are sorted according to the average price per unit of good. Hence, we get  $v(j) \geq \sum_{b \in B^o} a(b) \cdot e(b)_j$ . Given that the above holds for every good  $j$ , this implies that  $\sum_{1 \leq j \leq m} v(j) \geq \sum_{b \in B^o} p(b)$ , as requested.

The above theorem is the central tool for proving the following theorem:

**Theorem 2** CAMUS is complete: it is guaranteed to find the optimal allocation in a multi-unit combinatorial auction problem.

## Pre-Processing of Singletons

In this subsection we explain the construction of the  $singleton_g$  vector described above, and demonstrate that  $singleton_g(j)$  is the revenue-maximizing set of singleton bids for good  $g$  that request a total not exceeding  $j$  units.

Let  $b_1, b_2, \dots, b_l$  be bids for a single good  $g$ , where the total number of available units of good  $g$  is  $q$ . Let  $p(b_i)$  and  $e(b_i)$  be the price offer and the quantity requested by  $b_i$ , respectively. Our aim is to compute the optimal selection of  $b_i$ 's in order to allocate  $k$  units of good  $g$ , for  $1 \leq k \leq q$ . Consider a two dimensional grid of size  $[1 \dots l] \times [1 \dots q]$  where the  $(i, j)$ -th entry, denoted by  $U(i, j)$ , is the optimal allocation of  $j$  units considering only bids  $b_1, b_2, \dots, b_i$ . The value of  $U(i, j)$ , denoted by  $V(i, j)$ , is the sum of the price offers of the bids in  $U(i, j)$ .  $U(1, j)$  will be  $b_1$  if  $b_1$  requests no more than  $j$  units, and otherwise will be the empty set. Now we can define  $U(i, j)$  recursively:

1.  $e(b_i) > j$ :  $U(i, j) = U(i-1, j)$ ;
2.  $e(b_i) = j$ : if  $p(b_i) > V(i-1, j)$  then  $U(i, j) = b_i$ . Else  $U(i, j) = U(i-1, j)$ .
3.  $e(b_i) < j$ : if  $V(i-1, j) \geq p(b_i) + V(i-1, j - e(b_i))$  then  $U(i, j) = U(i-1, j)$ . Else  $U(i, j) = b_i \cup U(i-1, j - e(b_i))$ .

This dynamic programming procedure is polynomial, and yields the desired result; the optimal allocation of  $k$  units is given by  $U(l, k)$ . Set  $singleton_g(k) = U(l, k)$ ,  $1 \leq k \leq q$ .

## Experimental results

Unfortunately, no real-world data exists to describe how bidders will behave in general multi-unit combinatorial auctions, precisely because the determination of winners in such auctions was previously unfeasible. We have therefore tested CAMUS on sets of bids drawn from a random distribution. We created bids as follows,

varying the parameters  $numgoods$  and  $numbids$ , and fixing the parameters  $units_{max} = 5$ ,  $avgprice_{base} = 50$ ,  $avgprice_{var} = 25$ ,  $prob_1 = 0.8$ ,  $prob_2 = 0.65$ ,  $price_{var} = 0.5$ :

1. Set the number of units that exist for each good:
  - (a) For each good  $i$ , randomly choose  $units_i$  from the range  $[1 \dots units_{max}]$ .
  - (b) If  $\sum_i units_i \neq \frac{numgoods \sum_{j=1}^{units_{max}} j}{units_{max}}$  (the expectation on  $\sum_i units_i$ ) then go to (a). This ensures that each trial involves the same total number of units.
2. Set an average price for each good:  $avgprice_i$  is drawn uniformly randomly from the range  $[avgprice_{base} - avgprice_{var} \dots avgprice_{base} + avgprice_{var}]$ .
3. Select the number of goods in the bid. This number is drawn from a decay distribution:
  - (a) Randomly choose a good that has not already been added to this bid
  - (b) With probability  $prob_1$ , if more goods remain then go to (a)
4. Select the number of units of each good, according to another decay distribution:
  - (a) Add a unit
  - (b) With probability  $prob_2$ , if more units remain then go to (a)
5. Set a price for this bid:  $price = rand(1 - price_{var}, 1 + price_{var}) \cdot \sum_{i \in bid} (avgprice_i \cdot units_i)$

This distribution has the following characteristics that we consider to be reasonable. Bids will tend to request a small number of goods, independent of the total number of goods. Such data cases are computationally harder than drawing a number of goods uniformly from a range, or than scaling the average number of goods per bid to the maximum number of goods. Likewise, bids will tend to name a small number of units per good. Prices tend to increase linearly in the number of units, for a fixed set of goods. This is a harder case for our pruning technique, much harder than drawing prices uniformly from a range. In fact, it may be reasonable for prices to be superlinear in the number of units, as the motivation for holding a CA in the first place may be that bidders are expected to value bundles more than individual goods. However, this would be an easier case for our pruning algorithm, so we tested on the linear case instead. The construction of realistic, hard data distributions remains a topic for further research.

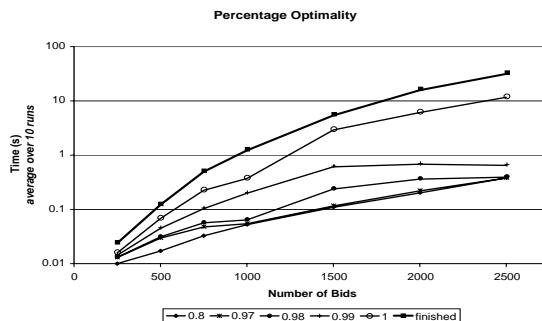
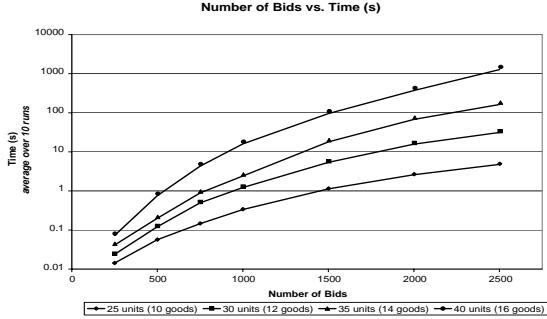
Our experimental data was collected on a Pentium III-733 running Windows 2000, with 25 MB allocated for CAMUS's cache. Our figure *Number of Bids vs Time* shows CAMUS's performance on the distribution described above, with each line representing runs with a different number of goods. Note that, for example, CAMUS solved problems with 35 objects (14 goods) and 2500 bids in about two minutes, and problems with 25 objects (10 goods) and 1500 bids in about a second. Because the lines in this graph are sub-linear on the logarithmic scale, CAMUS's performance is sub-exponential in the number of bids, though it remains exponential in the number of goods. Our figure *Percentage Optimality* shows CAMUS's anytime performance. Each line on the graph shows the time taken to find solutions with revenue that is some percentage of the optimal, calculated after the algorithm terminated. Note that the time taken to *find* the optimal solution is less than the time taken for the algorithm to finish, *proving* that this solution is optimal. These

anytime results are very encouraging—note that CAMUS finds a 99% optimal solution an order of magnitude more quickly than it takes for the algorithm to run to completion. This suggests that CAMUS could be useful on much larger problems than we have shown here if an optimal solution were not required.

## Conclusions

In this paper we introduced CAMUS, a novel algorithm for determining the optimal set of winning bids in general multi-unit combinatorial auctions. The algorithm has been tested on a variety of data distributions and has been found to solve problems of considerable scale in an efficient manner. CAMUS extends our CASS algorithm for single-unit combinatorial auctions, and enables a wide extension of the class of combinatorial auctions that can be efficiently implemented. In our current research we are studying the addition of random noise into our good and bin ordering heuristics, combined with periodic restarts and the deletion of previously-searched bids, to improve performance on hard cases while still retaining completeness.

- Monderer, D., and Tennenholtz, M. 2000. Optimal Auctions Revisited. *Artificial Intelligence*, forthcoming.
- Nisan, N., and Ronen, A. 2000. Computationally feasible vcg mechanisms. To appear.
- Nisan, N. 1999. Bidding and allocation in combinatorial auctions. Working paper.
- Rothkopf, M.; Pekec, A.; and Harstad, R. 1998. Computationally manageable combinatorial auctions. *Management Science* 44(8):1131–1147.
- Sandholm, T. 1999. An algorithm for optimal winner determination in combinatorial auctions. In *IJCAI-99*.
- Tennenholtz, M. 2000. Some tractable combinatorial auctions. To appear in the proceedings of AAAI-2000.
- Vries, S., and Vohra, R. 2000. Combinatorial auctions: A brief survey. Unpublished manuscript.
- Wellman, M.; Wurman, P.; Walsh, W.; and MacKie-Mason, J. 1998. Auction protocols for distributed scheduling. Working paper (to appear in Games and Economic Behavior).



## References

- Fudenberg, D., and Tirole, J. 1991. *Game Theory*. MIT Press.
- Fujishima, Y.; Leyton-Brown, K.; and Shoham, Y. 1999. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *IJCAI-99*.
- Lehmann, D.; O'Callaghan, L.; and Shoham, Y. 1999. Truth revelation in rapid, approximately efficient combinatorial auctions. In *ACM Conference on Electronic Commerce*.

# Towards a Universal Test Suite for Combinatorial Auction Algorithms

Kevin Leyton-Brown  
 Dept. of Computer Science  
 Stanford University  
 Stanford, CA 94305  
 kevinlb@stanford.edu

Mark Pearson  
 Dept. of Computer Science  
 Stanford University  
 Stanford, CA 94305  
 mpearson@stanford.edu

Yoav Shoham  
 Dept. of Computer Science  
 Stanford University  
 Stanford, CA 94305  
 shoham@cs.stanford.edu

## ABSTRACT

General combinatorial auctions—auctions in which bidders place unrestricted bids for bundles of goods—are the subject of increasing study. Much of this work has focused on algorithms for finding an optimal or approximately optimal set of winning bids. Comparatively little attention has been paid to methodical evaluation and comparison of these algorithms. In particular, there has not been a systematic discussion of appropriate data sets that can serve as universally accepted and well motivated benchmarks. In this paper we present a suite of distribution families for generating realistic, economically motivated combinatorial bids in five broad real-world domains. We hope that this work will yield many comments, criticisms and extensions, bringing the community closer to a universal combinatorial auction test suite.<sup>1</sup>

## 1. INTRODUCTION

### 1.1 Combinatorial Auctions

Auctions are a popular way to allocate goods when the amount that bidders are willing to pay is either unknown or unpredictably changeable over time. The rise of electronic commerce has facilitated the use of increasingly complex auction mechanisms, making it possible for auctions to be applied to domains for which the more familiar mechanisms are inadequate. One such example is provided by combinatorial auctions (CA's), multi-object auctions in which bids name bundles of goods. These auctions are attractive because they allow bidders to express *complementarity* and *substitutability* relationships in their valuations for sets of goods. Because CA's allow bids for arbitrary bundles of goods, an agent may offer a different price for some bundle of goods than he offers for the sum of his bids for its disjoint

subsets; in the extreme case he may bid for a bundle with the guarantee that he will not receive any of its subsets. An example of complementarity is an auction of used electronic equipment, in which a bidder values a particular TV at  $x$  and a particular VCR at  $y$  but values the pair at  $z > x + y$ . An agent with substitutable valuations for two copies of the same book might value either single copy at  $x$ , but value the bundle at  $z < 2x$ . In the special case where  $z = x$  (the agent values a second book at 0, having already bought a first) the agent may submit the set of bids  $\{bid_1 \text{ XOR } bid_2\}$ . By default, we assume that any satisfiable sets of bids that are not explicitly XOR'ed is a candidate for allocation. We call an auction in which all goods are distinguishable from each other a *single-unit CA*. In contrast, in a multi-unit CA some of the goods are indistinguishable (e.g., many identical TVs and VCRs) and bidders request some number of goods from each indistinguishable set. This paper is primarily concerned with single-unit CA's, since most research to date has been focused on this problem. However, when appropriate we will discuss ways that our distributions could be generalized to apply to multi-unit CA's.

### 1.2 The Computational Combinatorial Auction Problem

In a combinatorial auction, a seller is faced with a set of price offers for various bundles of goods, and his aim is to allocate the goods in a way that maximizes his revenue. (For an overview of this problem, see [8].) This optimization problem is intractable in the general case, even when each good has only a single unit. Because of the intractability of general CA's, much research has focused on subcases of the CA problem that are tractable; see [22] and more recently [25]. However, these subcases are very restrictive and therefore are not applicable to many CA domains. Other research attempts to define mechanisms within which general CA's will be tractable (achieved by various trade-offs including bid withdrawal penalties, activity rules and possible inefficiency). Milgrom [15] defines the Simultaneous Ascending Auction mechanism which has been very influential, particularly in the recent FCC spectrum auctions. However, this approach has drawbacks, discussed for example in [6]. In the general case there is no substitute for a completely unrestricted CA. Consequently, many researchers have recently begun to propose algorithms for determining the winners of a general CA, with encouraging results. This wave of research has given rise to a new problem, however. In order to test (and thus to improve) such algorithms, it has been

<sup>1</sup>This work was partly supported by DARPA grant number F30602-98-C-0214-P00005.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'00, October 17-20, 2000, Minneapolis, Minnesota.  
 Copyright 2000 ACM 1-58113-272-7/00/0010 ..\$5.00

necessary to use some sort of test suite. Since general CA's have never been widely held, there is no data recording the bidding behavior of real bidders upon which such a test suite may be built. In the absence of such natural data, we are left only with the option of generating artificial data that is representative of the sort of scenarios one is likely to encounter. The goal of this paper is to facilitate the creation of such a test suite.

## 2. PAST WORK ON TESTING CA ALGORITHMS

### 2.1 Experiments with Human Subjects

One approach to experimental work on combinatorial auctions uses human subjects. These experiments assign valuation functions to subjects, then have them participate in auctions using various mechanisms [3, 12, 7]. Such tests can be useful for understanding how real people bid under different auction mechanisms; however, they are less suitable for evaluating the mechanisms' computational characteristics. In particular, this sort of test is only as good as the subjects' valuation functions, which in the above papers were hand-crafted. As a result, this technique does not easily permit arbitrary scaling of the problem size, a feature that is important for characterizing an algorithm's performance. In addition, this method relies on relatively naive subjects to behave rationally given their valuation functions, which may be unreasonable when subjects are faced with complex and unfamiliar mechanisms.

### 2.2 Particular Problems

A parallel line of research has examined particular problems to which CA's seem well suited. For example, researchers have considered auctions for the right to use railroad tracks [5], real estate [19], pollution rights [13], airport time slot allocation [21] and distributed scheduling of machine time [26]. Most of these papers do not suggest holding an unrestricted general CA, presumably because of the computational obstacles. Instead, they tend to discuss alternative mechanisms that are tailored to the particular problem. None of them proposes a method of generating test data, nor does any of them describe how the problem's difficulty scales with the number of bids and goods. However, they still remain useful to researchers interested in general CA's because they give specific descriptions of problem domains to which CA's may be applied.

### 2.3 Artificial Distributions

Recently, a number of researchers have proposed algorithms for determining the winners of general CA's. In the absence of test suites, some suggested novel bid generation techniques, parameterized by number of bids and goods [24, 10, 4, 8]. (Other researchers have used one or more of these distributions, e.g., [17], while still others have refrained from testing their algorithms altogether, e.g., [16, 14].) Parameterization represents a step forward, making it possible to describe performance with respect to the problem size. However, there are several ways in which each of these bid generation techniques falls short of realism, concerning the selection of which goods and how many goods to request in a bundle, what price to offer for the bundle, and which bids to combine in an XOR'ed set. More fundamentally, however, all of these approaches suffer from failing to

model bidders explicitly, and from attempting to represent an economic situation with a non-economic model.

#### 2.3.1 Which goods

First, each of the distributions for generating test data discussed above has the property that all bundles of the same size are equally likely to be requested. This assumption is clearly violated in almost any real-world auction: most of the time, certain goods will be more likely to appear together than others. (Continuing our electronics example, TVs and VCRs will be requested together more often than TVs and printers.)

#### 2.3.2 Number of goods

Likewise, each of the distributions for generating test data determines the number of goods in a bundle completely independently from determining *which* goods appear in the bundle. While this assumption appears more reasonable it will still be violated in many domains, where the expected length of a bundle will be related to which goods it contains. (For example, people buying computers will tend to make long combinatorial bids, requesting monitors, printers, etc., while people buying refrigerators will tend to make short bids.)

#### 2.3.3 Price

Next, there are problems with the pricing<sup>2</sup> schemes used by all four techniques. Pricing is especially crucial: if prices are not chosen carefully then an otherwise hard distribution can become computationally easy.

In Sandholm [24] prices are drawn randomly from either  $[0, 1]$  or from  $[0, g]$ , where  $g$  is the number of goods requested. The first method is clearly unreasonable (and computationally trivial) since price is unrelated to the number of goods in a bid—note that a bid for many goods and for a small subset of the same bid will have exactly the same price on expectation. The second is better, but has the disadvantage that average and range are parameterized by the same variable.

In Boutilier et al.[4] prices of bids are distributed normally with mean 16 and standard deviation 3, giving rise to the same problem as the  $[0, 1]$  case above.

In Fujishima et al.[10] prices are drawn from  $[g(1-d), g(1+d)]$ ,  $d = 0.5$ . While this scheme avoids the problems described above, prices are simply additive in  $g$  and are unrelated to *which* goods are requested in a bundle, both unrealistic assumptions in some domains.

More fundamentally, Andersson et al.[1] note a critical pricing problem that arises in several of the schemes discussed above. As the number of bids to be generated becomes large, a given short bid will be drawn much more frequently than a given long bid. Since the highest-priced bid for a bundle dominates all other bids for the same bundle, short bids end up being much more competitive. Indeed, it is pointed out that for extremely large numbers of bids a good approximation to the optimal solution is simply to take the best singleton bid for each good. One solution to this problem is to guarantee that a bid will

---

<sup>2</sup>Most of the existing literature on artificial distributions in combinatorial auctions refers to the monetary amount associated with a bundle as a "price". In Section 3 we will advocate the use of different terminology, but in this section we use the existing term for clarity.

be placed for each bundle at most once (for example, this approach is taken by Sandholm[24]). However, this solution has the drawback that it is unrealistic: different real bidders *are* likely to place bids on some of the same bundles.

Another solution to this problem is to make bundle prices superadditive in the number of goods they request—an assumption that may also be reasonable in many CA domains. A similar approach is taken by DeVries and Vohra [8], who make the price for a bid a quadratic function of the prices of bids for subsets. For some domains this pricing scheme may result in too large an increase in price as a function of bundle length. The distributions presented in this paper will include a pricing scheme that may be configured to be superadditive or subadditive in bundle length, where appropriate, parameterized to control how rapidly the price offered increases or decreases as a function of bundle length.

#### 2.3.4 XOR bids

Finally, while most of the bid-generation techniques discussed above permit bidders to submit sets of bids XOR’ed together, they have no way of generating meaningful sets of such bids. As a consequence the computational impact of XOR’ed bids has been very difficult to characterize.

### 3. GENERATING REALISTIC BIDS

While the lack of standardized, realistic test cases does not make it impossible to evaluate or compare algorithms, it does make it difficult to know what magnitude of real-world problems each algorithm is capable of solving, or what features of real-world problems each algorithm is capable of exploiting. This second ambiguity is particularly troubling: it is likely that algorithms would be designed *differently* if they took the features of more realistic<sup>3</sup> bidding into account.

#### 3.1 Prices, price offers and valuations

The term “price” has traditionally been used by researchers constructing artificial distributions to describe the amount offered for a bundle. However, this term really refers to the amount a bidder is made to pay for a bundle, which is of course mechanism-specific and is often not the same as the amount offered. Indeed, it is impossible to model bidders’ price offers at all without committing to a particular auction mechanism. In the distributions described in this paper, we will assume a sealed-bid incentive-compatible mechanism, where the price offered for a bundle is equal to the bidder’s valuation. Hence, in the rest of this paper, we will use the terms *price offer* and *valuation* interchangeably. Researchers wanting to model bidding behavior in other mechanisms could transform the valuation generated by our distributions according to bidders’ equilibrium strategies in the new mechanism.

#### 3.2 The CATS suite

---

<sup>3</sup>Previous work characterizes hard cases for weighted set packing—equivalent to the combinatorial auction problem. Real-world bidding is likely to exhibit various regularities, however, as discussed throughout this paper. A data set designed to include the same regularities may be more useful for predicting the performance of an algorithm in a real-world auction.

In this paper we present CATS (Combinatorial Auction Test Suite), a suite of distributions for modeling realistic bidding behavior. This suite is grounded in previous research on specific applications of combinatorial auctions, as described in section 2.1 above. At the same time, all of our distributions are parameterized by number of goods and bids, facilitating the study of algorithm performance. This suite represents a move beyond current work on modeling bidding in combinatorial auctions because we provide an economic motivation for both the contents and the valuation of a bundle, deriving them from basic bidder preferences. In particular, in each of our distributions:

- Certain goods are more likely to appear together than others.
- The number of goods appearing in the bundle is often related to which goods appear in the bundle.
- Valuations are related to which goods appear in the bundle. Where appropriate, valuations can be configured to be subadditive, additive or superadditive in the number of goods requested.
- Sets of XOR’ed bids are constructed in meaningful ways, on a per-bidder basis.

We do not intend for this paper to stand as an isolated statement on bidding in combinatorial auctions, but rather as the beginning of a dialogue. We hope to receive many suggestions and criticisms from members of the CA community, enabling us both to update the distributions proposed here and to include distributions modeling new domains. In particular, our distributions include many parameters, for which we suggest default values. Although these values have evolved somewhat during our development of the test suite, it has not yet been possible to understand the role each parameter plays in the difficulty or realism of the resulting distribution, and our choice may be seen as highly subjective. We hope and expect to receive criticisms about these parameter values; for this reason we include a CATS version number with the defaults to differentiate them from future defaults. The suite also contains a legacy section including all bid generation techniques described above, so that new algorithms may easily be compared to previously-published results. More information on our test suite, including executable versions of our distributions for Solaris, Linux and Windows may be found at <http://robotics.stanford.edu/CATS>.

In section 4, below, we present distributions based on five real-world situations. For most of our distributions, the mechanism for generating bids requires first building a graph representing adjacency relationships between goods. Later, the mechanism uses the graph, generated in an economically-motivated way, to derive complementarity properties between goods and substitutability properties for bids. Of the five real-world situations we model, the first three concern complementarity based on adjacency in (physical or conceptual) space, while the final two concern complementarity based on correlation in time. Our first example (4.1) models shipping, rail and bandwidth auctions. Goods are represented as edges in a nearly planar graph, with agents submitting an XOR’ed set of bids for paths connecting two nodes. Our second example (4.2) models an auction of real estate, or more generally of any goods over which two-dimensional

adjacency is the basis of complementarity. Again the relationship between goods is represented by a graph, in this case strictly planar. In (4.3) we relax the planarity assumption from the previous example in order to model arbitrary complementarities between discrete goods such as electronics parts or collectables. Our fourth example (4.4) concerns the matching of time-slots for a fixed number of different goods; this case applies to airline take-off and landing rights auctions. In (4.5) we discuss the generation of bids for a distributed job-shop scheduling domain, and also its application to power generation auctions. Finally, in (4.6), we provide a legacy suite of bid generation techniques, including all those discussed in (2.3) above.

In the description of the distributions that follow, let  $\text{rand}(a, b)$  represent a real number drawn uniformly from  $[a, b]$ . Let  $\text{rand\_int}(a, b)$  represent a random integer drawn uniformly from the same interval. With respect to a given graph, let  $e(x, y)$  represent the proposition that an edge exists between nodes  $x$  and  $y$ . Denote the number of goods in a bundle  $B$  as  $|B|$ . The statement a good  $g$  is in a bundle  $B$  means that  $g \in B$ . All of the distributions presented here are parameterized by the number of goods (*num\_goods*) and number of bids (*num\_bids*).

## 4. CATS IN DETAIL

### 4.1 Paths in Space

There are many real-world problems involving bidding on paths in space. Generally, this class may be characterized as the problem of purchasing a connection between two points. Examples include truck routes [23], natural gas pipeline networks [20], network bandwidth allocation, and the right to use railway tracks [5].<sup>4</sup> In particular, spatial path problems consist of a set of points and accessibility relations between them. Although the distribution we propose may be configured to model bidding in any of the above domains, we will use the railway domain as our motivating example since it is both intuitive and well-understood.

More formally, we will represent this railroad auction by a graph in which each node represents a location on a plane, and an edge represents a connection between locations. The goods at auction are therefore the edges of the graph, and bids request a set of edges that form a path between two nodes. We assume that no bidder will desire more than one path connecting the same two nodes, although the bidder may value each path differently.

#### 4.1.1 Building the Graph

The first step in modeling bidding behavior for this problem is determining the graph of spatial and connective relationships between cities. One approach would be to use an actual railroad map, which has the advantage that the resulting graph would be unarguably realistic. However,

<sup>4</sup>Electric power distribution is a frequently discussed real world problem which seems superficially similar to the problems discussed here. However, many of the complementarities in this domain arise from physical laws governing power flow in a network. Consideration of these laws becomes very complex in networks of interesting size. Also, because these laws are taken into account during the construction of power networks, the networks themselves are difficult to model using randomly generated graphs. For these reasons, we do not attempt to model this domain.

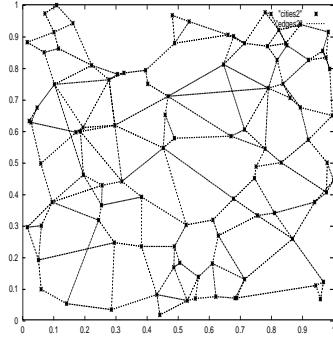


Figure 1: Sample Railroad Graph

it would be difficult to find a set of real-world maps that could be said to exhibit a similar sort of connectivity and would encompass substantial variation in the number of cities. Since scalability of input data is of great importance to the testing of new CA algorithms, we have chosen to propose generating such graphs randomly. Our technique for generating graphs has various parameters that may be adjusted as necessary; in our opinion it produces realistic graphs with the recommended settings. Figure 1 shows a representative example of a graph generated using our technique.

We begin with *num\_cities* nodes randomly placed on a plane. We add edges to this graph,  $G$ , starting by connecting each node to a fixed number of its nearest neighbors. Next, we iteratively consider random pairs of nodes and examine the shortest path connecting them, if any. To compare, we also compute various alternative paths that would require one or more edges to be added to the graph, given a penalty proportional to distance for adding new edges. (We do this by considering a complete graph  $C$ , an augmentation of  $G$  with new edges weighted to reflect the distance penalty.) If the shortest path involves new edges—despite the penalty—then the new edges (without penalty) are added to  $G$ , and replace the existing edges in  $C$ . This process models our simplifying assumption that there will exist uniform demand for shipping between any pair of cities, though of course it does not mimic the way new links would actually be added to a rail network. Our technique produces slightly non-planar graphs—graphs on a plane in which edges occasionally cross at points other than nodes. We consider this to be reasonable, as the same phenomenon may be observed in real-world rail lines, highways, network wiring, etc. Determining the “reasonableness” of a graph is of course a subjective task unless more quantitative metrics are used to assess quality; we see the identification and application of such metrics (for this and other distributions) as an important topic for future work.

#### 4.1.2 Generating Bids

Given a map of cities and the connectivity between them, there is the orthogonal problem of modeling bidding itself. We propose a method which generates a set of substitutable bids from a hypothetical agent’s point of view. We start with the value to an agent for shipping from one city to another and with a shipping cost which we make equal to the Euclidean distance between the cities. We then place XOR bids on all paths on which the agent would make a profit

```

Let  $num\_cities = f(num\_goods)$ 
Randomly place nodes (cities) on a unit box
Connect each node to its initial_connections
nearest neighbors
For  $i = 1$  to  $num\_building\_paths$ :
   $C = G$ 
  For every pair of nodes  $n_1, n_2 \in G$  where
     $\neg e(n_1, n_2)$ :
    Add an edge to  $C$  of length
       $building\_penalty \cdot Euclidean\_distance(n_1, n_2)$ 
  Choose two nodes at random, and find the
  shortest path between them in  $C$ 
  If shortest path uses edges that do not
  exist in  $G$ :
    For every such pair of nodes
       $n_1, n_2 \in G$  add an edge to  $G$  with
      length  $Euclidean\_distance(n_1, n_2)$ 
  End If
End For
If total number of edges in  $G \neq num\_goods$ ,
  restart

```

**Figure 2: Graph-Building Technique**

```

While  $num\_generated\_bids < num\_bids$ :
  Randomly choose two nodes,  $n_1$  and  $n_2$ 
   $d = rand(1, shipping\_cost\_factor)$ 
   $cost = Euclidean\_distance(city_1, city_2)$ 
   $value = d \cdot Euclidean\_distance(city_1, city_2)$ 
  Make XOR bids of  $value - cost$  on every path
    from  $city_1$  to  $city_2$  with  $cost < value$ 
  If there are more than  $max\_bid\_set\_size$  such
    paths, bid on the  $max\_bid\_set\_size$  paths
    that maximize  $value - cost$ .
End While

```

**Figure 3: Bid-Generation Technique**

(i.e., those paths where  $utility - cost > 0$ ). The path's value is random, in (parameterized) proportion to the Euclidean distance between the chosen cities. Since the shipping cost is the Euclidean distance between two cities, we use this as the lower bound for value as well, since only bidders with such valuations would actually place bids.

Note that this distribution, and indeed all others presented in this paper, may generate slightly more than  $num\_bids$  bids. In our experience CA optimization algorithms tend not to be highly sensitive in the number of bids, so we judged it more important to build economically sensible sets of substitutable bids. When generating a precise number of bids is important, an appropriate number of bids may be removed after all bids have been generated so that the total will be met exactly.

Note that 1 is used as a lower bound for  $d$  because any bidder with  $d < 1$  would find no profitable paths and therefore would not bid.

This is CATS 1.0 problem 1. CATS default parameters:  $initial\_connections = 2$ ,  $building\_penalty = 1.7$ ,  $num\_building\_paths = num\_cities^2/4$ ,  $shipping\_cost\_factor = 1.5$ ,  $max\_bid\_set\_size = 5$ , and  $f(num\_goods) = 0.529689 * NUMGOODS + 3.4329$ .

#### 4.1.3 Multi-Unit Extensions: Bandwidth Allocation, Commodity Flow

This model may also be used to generate realistic data

```

Place nodes at integer vertices  $(i, j)$  in a
plane, where  $1 \leq i, j \leq \lceil \sqrt{(num\_goods)} \rceil$ 
For each node  $n$ :
  If  $n$  is on the edge of the map
    Connect  $n$  to as many hv-neighbors as
    possible
  Else
    If  $rand(0, 1) \leq three\_prob$ 
      Connect  $n$  to a random set of
      three of its four hv-neighbors
    Else
      Connect  $n$  to all four of its
      hv-neighbors
    While  $rand(0, 1) \leq additional\_neighbor$ :
      Connect  $g$  to one of its
      d-neighbors, provided that the
      new diagonal edge will not
      cross another diagonal edge
  End While
End For

```

**Figure 4: Graph-Building Technique**

for multi-unit CA problems such as network bandwidth allocation and general commodity flow. The graph may be created as above, but with a number of units (capacity) assigned to each edge. Likewise, the bidding technique remains unchanged except for the assignment of a number of units to each bid.

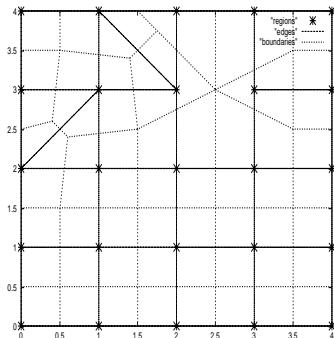
## 4.2 Proximity in Space

There is a second broad class of real-world problems in which complementarity arises from adjacency in two-dimensional space. An intuitive example is the sale of adjacent pieces of real estate [19]. Another example is drilling rights, where it is much cheaper for an (e.g.) oil company to drill in adjacent lots than in lots that are far from each other. In this section, we first propose a graph-generation mechanism that builds a model of adjacency between goods, and then describe a technique for generating realistic bids on these goods. Note that in this section nodes of the graph represent the goods at auction, while edges represent the adjacency relationship.

### 4.2.1 Building the Graph

There are a number of ways we could build an adjacency graph. The simplest would be to place all the goods (locations, nodes) in a grid, and connect each to its four neighbors. We propose a slightly more complex method in order to permit a variable number of neighbors per node (equivalent to non-rectangular pieces of real estate). As above we place all goods on a grid, but with some probability we omit a connection between goods that would otherwise represent vertical or horizontal adjacency, and with some probability we introduce a connection representing diagonal adjacency. (We call horizontally- or vertically-adjacent nodes *hv-neighbors* and diagonally-adjacent nodes *d-neighbors*.)

Figure 5 shows a sample real estate graph, generated by the technique described in Figure 4. Nodes of the graph are shown with asterisks, while edges are represented by solid lines. The dashed lines show one set of property boundaries that would be represented by this graph. Note that one node falls inside each piece of property, and that two pieces of property border each other iff their nodes share an edge.



**Figure 5:** Sample Real Estate Graph

#### 4.2.2 Generating Bids

To model realistic bidding behavior, we generate a set of common values for each good, and private values for each good for each bidder. The common value represents the appraised or expected resale value of each individual good. The private value represents how much one particular bidder values that good, as an offset to the common value (e.g., a private value of 0 for a good represents agreement with the common value). These private valuations describe a bidder's preferences, and so they are used to determine both a value for a given bid and the likelihood that a bidder will request a bundle that includes that good. There are two additional components to each bidder's preferences: a minimum total common value, and a budget. The former reflects the idea that a bidder may only wish to acquire goods of a certain recognized value. The latter reflects the fact that a bidder may not be able to afford every bundle that is of interest to him.

To generate bids, we first add a random good, weighted by a bidder's preferences, to the bidder's bid. Next, we determine whether another good should be added by drawing a value uniformly from  $[0,1]$ , and adding another good if this value is smaller than a threshold. This is equivalent to drawing the number of goods in a bid from a decay distribution.<sup>56</sup> We must now decide which good to add. First we allow a small chance that a new good will be added uniformly at random from the set of goods, without the requirement that it be adjacent to a good in the current bundle  $B$ . (This permits bundles requesting unconnected regions of the graph: for example, a hotel company may only wish to build in a city if it can acquire land for two hotels on opposite sides of the city.) Otherwise, we select a good from the set of nodes bordering the goods in  $B$ . The probability that some adjacent good

<sup>5</sup>We use Sandholm's [24] term "decay" here, though the distribution goes by various names—for a description of the distribution please see Section 4.6.1.

<sup>6</sup>There are two reasons we use a decay distribution here. First, we expect that most bids will request small bundles; a uniform distribution, on the other hand, would be expected to have the same number of bids for bundles of each cardinality. Also, bids for large bundles will often be computationally easier for CA algorithms than bids for small bundles, because choosing the former more highly restricts the future search. Second, we require a distribution where the expected bundle size is unaffected by changes in the total number of goods. Some other distributions, such as uniform and binomial, do not have this property.

```
Routine Add_Good_to_Bundle(bundle B)
  If rand(0, 1) ≤ jump_prob:
    Add a good  $g \notin B$  to  $B$ , chosen
      uniformly at random
  Else:
    Compute  $s = \sum_{x \notin B, y \in B, e(x,y)} pn(x)$  [pn()
      is defined below]
    Choose a random node  $x \notin B$  from the
      distribution  $\sum_{y \in B, e(x,y)} \frac{pn(x)}{s}$ 
    Add  $x$  to  $B$ 
  End If
End Routine
```

**Figure 6:** Add\_Good\_to\_Bundle for Spatial Proximity

$n_1$  will be added depends on how many edges  $n_1$  shares with the current bundle, and on the bidder's relative private valuations for  $n_1$  and  $n_2$ . For example, if nodes  $n_1$  and  $n_2$  are each connected to  $B$  by one edge, and the private valuation for  $n_1$  is twice that for  $n_2$  then the probability of adding  $n_1$  to  $B$ ,  $p(n_1)$ , is  $2p(n_2)$ . Further, if  $n_1$  has 3 edges to nodes in  $B$  while  $n_2$  is connected to  $B$  by only 1 edge, and the goods have equivalent private values, then  $p(n_1) = 3p(n_2)$ . Once we have determined all the goods in a bundle we set the price offered for the bundle, which depends on the sum of common and private valuations for the goods in the bundle, and also includes a function that is superadditive (with our parameter settings) in the number of goods.<sup>7</sup> Finally, we generate additional bids that are substitutable for the original bid, with the constraint that each bid in the set requests at least one good from the original bid.

This is CATS 1.0 problem 2. CATS default parameters:  $three\_prob = 1.0$ ,  $additional\_neighbor = 0.2$ ,  $max\_good\_value = 100$ ,  $max\_substitutable\_bids = 5$ ,  $additional\_location = 0.9$ ,  $jump\_prob = 0.05$ ,  $additivity = 0.2$ ,  $deviation = 0.5$ ,  $budget\_factor = 1.5$ ,  $resale\_factor = 0.5$ , and  $S(n) = n^{1+additivity}$ . Note that  $additivity = 0$  gives additive bids, and  $additivity < 0$  gives sub-additive bids.

#### 4.2.3 Spectrum Auctions

A related problem is the auction of radio spectrum, in which a government sells the right to use specific segments of spectrum in different geographical areas[18, 2].<sup>8</sup> It is possible to approximate bidding behavior in spectrum auctions by making the assumption that all complementarity arises from spatial proximity.<sup>9</sup> In this case, our spatial proximity model can also be used to generate realistic bidding distributions for spectrum auctions. The main difference between this problem and the real estate problem is that in a spectrum auction each good may have multiple units (frequency bands) for sale. It is insufficient to model this as a multi-unit CA problem, however, if bidders have the constraint

<sup>7</sup>Recall the discussion in Section 2.3.3 motivating the use of superadditive valuations.

<sup>8</sup>Spectrum auctions have not historically been formulated as general CA's, but the possibility of doing so is now being explored.

<sup>9</sup>This assumption would be violated, for example, if some bidders wanted to secure some spectrum in all metropolitan areas. Clearly the problem of realistic test data for spectrum auctions remains an area for future work.

```

For all  $g, c(g) = \text{rand}(1, \text{max\_good\_value})$ 
While  $\text{num\_generated\_bids} < \text{num\_bids}$ :
    For each good, reset
         $p(g) = \text{rand}(-\text{deviation} \cdot \text{max\_good\_value}, \text{max\_good\_value})$ 
         $pn(g) = \frac{p(g) + \text{deviation} \cdot \text{max\_good\_value}}{2 \cdot \text{deviation} \cdot \text{max\_good\_value}}$ 
    Normalize  $pn(g)$  so that  $\sum_g pn(g) = 1$ 
     $B = \{\}$ 
    Choose a node  $g$  at random, weighted by
         $pn()$ , and add it to  $B$ 
    While  $\text{rand}(0, 1) \leq \text{additional\_location}$ 
        Add_Good_to_Bundle( $B$ )
         $\text{value}(B) = \sum_{x \in B} (c(x) + p(x)) + S(|B|)$ 
    If  $\text{value}(B) \leq 0$  on  $B$ , restart bundle
        generation for this bidder
    Bid  $\text{value}(B)$  on  $B$ 
     $\text{budget} = \text{budget\_factor} \cdot \text{value}(B)$ 
     $\text{min\_resale\_value} = \text{resale\_factor} \cdot \sum_{x \in B} c(x)$ 
    Construct substitutable bids. For each
    good  $g_i \in B$ :
        Initialize a new bundle,  $B_i = \{g_i\}$ 
        While  $|B_i| < |B|$ :
            Add_Good_to_Bundle( $B_i$ )
            Compute  $c_i = \sum_{x \in B_i} c(x)$ 
    End For
    Make XOR bids on all  $B_i$  where
     $0 \leq \text{value}(B) \leq \text{budget}$  and
     $c_i \geq \text{min\_resale\_value}$ .
    If there are more than
         $\text{max\_substitutable\_bids}$  such bundles, bid
        on the  $\text{max\_substitutable\_bids}$  bundles
        having the largest value
End While

```

**Figure 7: Bid-Generation Technique**

that they want the same frequency in each region.<sup>10</sup> Instead, the problem can be modeled with multiple distinct goods per node in the graph, and bids constructed so that all nodes added to a bundle belong to the same ‘frequency’. With this method, it is also easy to incorporate other preferences, such as preferences for different types of goods. For instance, if two different types of frequency bands are being sold, one 5 megahertz wide and one 2.5 megahertz wide, an agent only wanting 5 megahertz bands could make substitutable bids for each such band in the set of regions desired (generating the bids so that the agent will acquire the same frequency in all the regions).

The scheme for generating price offers used in our real estate example may be inappropriate for the spectrum auction domain. Research indicates that while price offers will still tend to be superadditive, this superadditivity may be quadratic in the population of the region rather than exponential in the number of regions [2]. CATS includes a quadratic pricing option that may be used with this problem, in which the common value term above is used as a measure of population. Please see the CATS documentation for more details.

<sup>10</sup>To see why this cannot be modeled as a multi-unit CA, consider an auction for three regions with two units each, and three bidders each wanting one unit of two goods. In the optimal allocation,  $b_1$  gets 1 unit of  $g_1$  and 1 unit of  $g_2$ ,  $b_2$  gets 1 unit of  $g_2$  and 1 unit of  $g_3$ , and  $b_3$  gets 1 unit of  $g_3$  and 1 unit of  $g_1$ . In this example there is no way of assigning frequencies to the units so that each bidder gets the same frequency in both regions.

```

Build a fully-connected graph with one node for
each good
Label each edge from  $n_1$  to  $n_2$  with a weight
 $d(n_1, n_2) = \text{rand}(0, 1)$ 

```

**Figure 8: Graph-Building Technique**

```

Routine Add_Good_to_Bundle(bundle  $B$ )
    Compute  $s = \sum_{x \notin B, y \in B} d(x, y) \cdot pn(x)$ 
    Choose a random node  $x \notin B$  from the
        distribution  $\sum_{y \in B} d(x, y) \cdot \frac{pn(x)}{s}$ 
    Add  $x$  to  $B$ 
End Routine

```

**Figure 9: Routine Add\_Good\_to\_Bundle for Arbitrary Relationships**

### 4.3 Arbitrary Relationships

Sometimes complementarities between goods will not be as universal as geographical adjacency, but some kind of regularity in the complementarity relationships between goods will still exist. Consider an auction of different, indivisible goods, e.g. for semiconductor parts or collectables, or for distinct multi-unit goods such as the right to emit some quantity of two different pollutants produced by the same industrial process. In this section we discuss a general way of modeling such arbitrary relationships.

#### 4.3.1 Building the Graph

We express the likelihood that a particular pair of goods will appear together in a bundle as being proportional to the weight of the appropriate edge of a fully-connected graph. That is, the weight of an edge between  $n_1$  and  $n_2$  is proportional to the probability that, having only  $n_1$  in our bundle, we will add  $n_2$ . Weights are only proportional to probabilities because we must normalize the sum of all weights from a given good to 1 in order to calculate a probability.

#### 4.3.2 Generating Bids

Our technique for modeling bidding is a generalization of the technique presented in the previous section. We choose a first good and then proceed to add goods one by one, with the probability of each new good being added depending on the current bundle. Note that, since in this section the graph is fully-connected, there is no need for the ‘jumping’ mechanism described above. The likelihood of adding a new good  $g$  to bundle  $B$  is proportional to  $\sum_{y \in B} d(x, y) \cdot p_i(x)$ . The first term  $d(x, y)$  represents the likelihood (independent of a particular bidder) that goods  $x$  and  $y$  will appear in a bundle together; the second,  $p_i(x)$ , represents bidder  $i$ ’s private valuation of the good  $x$ . We implement this new mechanism by changing the routine `Add_Good_to_Bundle()`. We are thus able to use the same techniques for assigning a value to a bundle, as well as for determining other bundles with which it is substitutable.

This is CATS 1.0 problem 3. CATS default parameters:  $\text{max\_good\_value} = 100$ ,  $\text{additional\_good} = 0.9$ ,  $\text{max\_substitutable\_bids} = 5$ ,  $\text{additivity} = 0.2$ ,  $\text{deviation} = 0.5$ ,  $\text{budget\_factor} = 1.5$ ,  $\text{resale\_factor} = 0.5$ , and  $S(n) = n^{1+\text{additivity}}$ .

#### 4.3.3 Multi-Unit Pollution Rights Auctions: Future Work

Bidding in pollution-rights auctions[18, 13] may be modeled through a multi-unit generalization of the technique presented in this section. In such auctions, the government sells companies the right to generate specific amounts of some pollutant. In the United States, though these auctions are widely used, sulfur-dioxide is the only chemical for which they are the primary method of control. Current US pollution-rights auctions may therefore be modeled as single good multi-unit auctions. If the government were to conduct pollution rights auctions for multiple pollutants in the future, however, bidding would be best-represented as a multi-unit ‘Arbitrary Complementarity’ problem. The problem belongs to this class because some sets of pollutants are more likely to be produced than others, yet the relationship between pollutants can not be modeled through any notion of adjacency. Should such auctions become viable in the future, we hope that a pollution-rights distribution will be added to CATS .

## 4.4 Temporal Matching

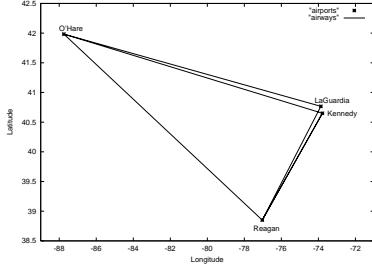
We now consider real-world domains in which complementarity arises from a temporal relationship between goods. In this section we discuss matching problems, in which corresponding time slices must be secured on multiple resources. The general form of temporal matching includes  $m$  sets of resources, in which each bidder wants 1 time slice from each of  $j \leq m$  sets subject to certain constraints on how the times may relate to one another (e.g., the time in set 2 must be at least two units later than the time in set 3). Here we concern ourselves with the problem in which  $j = 2$ , and model the problem of airport take-off and landing rights. Rassenti et al. [21] made the first study of auctions in this domain. The problem has been the topic for much other work; in particular [11] includes detailed experiments and an excellent characterization of bidder behavior.

The airport take-off and landing problem arises because certain high-traffic airports require airlines to purchase the right to take off or land during a given time slice. However, if an airline buys the right for a plane to take off at one airport then it must also purchase the right for the plane to land at its destination an appropriate amount of time later. Thus, complementarity exists between certain pairs of goods, where goods are the right to use the runway at a particular airport at a particular time. Substitutable bids are different departure/arrival packages; therefore bids will only be substitutable within certain limits.

### 4.4.1 Building the Graph

Departing from our graph-based approach above, we ground this example in the real map of high-traffic US airports for which the Federal Aviation Administration auctions take-off and landing rights, described in [11]. These are the four busiest airports in the United States: La Guardia International, Ronald Reagan Washington National, John F. Kennedy International, and O’Hare International. This map is shown below.

We chose not to use a random graph in this example because the number of bids and goods is dependent on the number of bidders and time slices at the given airports; it is not necessary to modify the number of *airports* in order to vary the problem size. Thus,  $\text{num\_cities} = 4$  and  $\text{num\_times} = [\text{num\_goods}/\text{num\_cities}]$ .



**Figure 10: Map of Airport Locations**

### 4.4.2 Generating Bids

Our bidding mechanism presumes that airlines have a certain tolerance for when a plane can take off and land (*early\_takeoff\_deviation*, *late\_takeoff\_deviation*, *early\_land\_deviation*, *late\_land\_deviation*), as related to their most preferred take-off and landing times (*start\_time*, *start\_time + min\_flight\_length*). We generate bids for all bundles that fit these criteria. The value of a bundle is derived from a particular agent’s utility function. We define a utility  $u_{max}$  for an agent, which corresponds to the utility the agent receives for flying from *city<sub>1</sub>* to *city<sub>2</sub>* if it receives the ideal takeoff and landing times. This utility depends on a common value for a time slot at the given airport, and deviates by a random amount. Next we construct a utility function which reduces  $u_{max}$  according to how late the plane will arrive, and how much the flight time deviates from optimal.

This is CATS 1.0 problem 4. CATS default parameters: *max\_airport\_value* = 5, *longest\_flight\_length* = 10, *deviation* = 0.5, *early\_takeoff\_deviation* = 1, *late\_takeoff\_deviation* = 2, *early\_land\_deviation* = 1, *late\_land\_deviation* = 2, *delay\_coeff* = 0.9, and *amount\_late\_coeff* = 0.75.

## 4.5 Temporal Scheduling

Wellman et al. [26] proposed distributed job-shop scheduling with one resource as a CA problem. We provide a distribution that mirrors this problem. While there exist many algorithms for solving job-shop scheduling problems, the distributed formulation of this problem places it in an economic context. In the problem formulation from Wellman et al., a factory conducts an auction for time-slices on some resource. Each bidder has a job requiring some amount of machine time, and a deadline by which the job must be completed. Some jobs may have additional, later deadlines which are less desirable to the bidder and so for which the bidder is willing to pay less.

### 4.5.1 Generating Bids

In the CA formulation of this problem, each good represents a specific time-slice. Two bids are substitutable if they constitute different possible schedules for the same job. We determine the number of deadlines for a given job according to a decay distribution, and then generate a set of substitutable bids satisfying the deadline constraints. Specifically, let the set of deadlines of a particular job be  $d_1 < \dots < d_n$  and the value of a job completed by  $d_1$  be  $v_1$ , superadditive in the job length. We define the value of a job completed by deadline  $d_i$  as  $v_i = v_1 \cdot \frac{d_1}{d_i}$ , reflecting the intuition that the

```

Set the average valuation for each city's
  airport: cost(city) = rand(0, max_airport_value)
Let max_l = length of longest distance between
  any two cities
While num_generated_bids < num_bids:

  Randomly select city1 and city2 where
    e(city1, city2)
    l = distance(city1, city2)
    min_flight_length =
      round(longest_flight_length ·  $\frac{1}{max_l}$ )
    start_time =
      rand_int(1, num_times - min_flight_length)
    dev = rand(1 - deviation, 1 + deviation)
    Make substitutable (XOR) bids. For
      takeoff =
        max(1, start_time - early_takeoff_deviation)
      to min(num_times, start_time +
        late_takeoff_deviation):
        For land = takeoff + min_flight_length
          to
            min(start_time + min_flight_length +
              late_land_deviation, num_times):
              amount_late =
                min(land - (start_time +
                  min_flight_length), 0)
              delay =
                land - takeoff - min_flight_length
              Bid dev · (cost(city1) + cost(city2)) ·
                delay_coeff ·
                amount_late_coeff · amount_late for
                takeoff at time takeoff at
                city1 and landing at time land
                at city2
        End For
    End For
End While

```

Figure 11: Bid-Generation Technique

decrease in value for a later deadline is proportional to its ‘lateness’.

Note that, like Wellman et al., we assume that all jobs are eligible to be started in the first time-slot. Our formulation of the problem differs in only one respect—we consider only allocations in which jobs receive continuous blocks of time. However, this constraint is not restrictive because for any arbitrary allocation of time slots to jobs there exists a new allocation in which each job receives a continuous block of time and no job finishes later than in the original allocation. (This may be achieved by numbering the winning bids in increasing order of scheduled end time, and then allocating continuous time-blocks to jobs in this order. Clearly no job will be rescheduled to finish later than its original scheduled time.) Note also that this problem cannot be translated to a trivial one-good multi-unit CA problem because jobs have different deadlines.

This is CATS 1.0 problem 5. CATS default parameters:  $deviation = 0.5$ ,  $prob\_additional\_deadline = 0.9$ ,  $additivity = 0.2$ , and  $max\_length = 10$ . Note that we propose a constant maximum job length, because the length of time a job requires should not depend on the amount of time the auctioneer makes available.

#### 4.5.2 Multi-Unit Power Generation Auctions: Future Work

```

While num_generated_bids < num_bids:

  l = rand_int(1, max_length)
  d1 = rand_int(l, num_goods)
  dev = rand(1 - deviation, 1 + deviation)
  cur_max_deadline = 0
  new_d = d1
  To generate substitutable (XOR) bids. Do:
    Make bids with price offered
      = dev · l + additivity · d1 / new_d for all
      blocks [start, end] where start ≥ 1,
      end ≤ new_d, end > cur_max_deadline,
      end - start = l
      cur_max_deadline = new_d
      new_d = rand_int(cur_max_deadline +
        1, num_goods)
    While rand(0, 1) ≤ prob_additional_deadline
  End While

```

Figure 12: Bid-Generation Technique

The problem of scheduling power generation is superficially similar to the job-shop scheduling problem described above. In these auctions, electrical power generation companies bid to produce a certain quantity of power for each hour of the day. This new problem differs from job-shop scheduling primarily because different kinds of power plants will exhibit very different utility functions, considering different sorts of goods to be complementary. For example, some plants will want to produce for long blocks of time (because they have startup and shutdown costs), others will prefer certain times of day due to labor costs, and still others will have neither restriction [9]. Due to the domain-specific complexity of bidder utilities, the construction of a distribution for this problem remains an area for future work.

## 4.6 Legacy Distributions

To aid researchers designing new CA algorithms by facilitating comparison with previous work, CATS includes the ability to generate bids according to all previous published test distributions of which we are aware, that are able to scale with the number of goods and bids. Each of these distributions may be seen as an answer to three questions: what number of goods to request in a bundle, which goods to request, and the price offered for a bundle. We begin by describing different techniques for answering each of these three questions, and then show how they have been combined in previously published work.

### 4.6.1 Number of Goods

**Uniform:** Uniformly distributed on  $[1, num\_goods]$

**Normal:** Normally distributed with  $\mu = \mu_{goods}$  and  $\sigma = \sigma_{goods}$

**Constant:** Fixed at  $constant\_goods$

**Decay:** Starting with 1, repeatedly increment the size of the bundle until  $rand(0, 1)$  exceeds  $\alpha$

**Binomial:** Request  $n$  goods with probability  $p^n(1-p)^{num\_goods-n} \binom{num\_goods}{n}$

**Exponential:** Request  $n$  goods with probability  $C \exp^{-n/q}$

### 4.6.2 Which Goods

**Random:** Draw  $n$  random goods from the set of all goods,

without replacement<sup>11</sup>

#### 4.6.3 Price Offer

**Fixed Random:** Uniform on [*low\_fixed*, *hi\_fixed*].

**Linear Random:** Uniform on [*low\_linearly*·*n*, *hi\_linearly*·*n*]

**Normal:** Draw from a normal distribution with  $\mu = \mu_{\text{price}}$  and  $\sigma = \sigma_{\text{price}}$

**Quadratic**<sup>12</sup>: For each good  $k$  and each bidder  $i$  set the value  $v_k^i = \text{rand}(0, 1)$ . Then  $i$ 's price offer for a set of goods  $S$  is  $\sum_{k \in S} v_k^i + \sum_{k, q} v_k^i v_q^i$ .

### 4.7 Previously Published Distributions

The following is a list of the distributions used in all published tests of which we are aware. In each case we describe first the method used to choose the number of goods, followed by the method used to choose the price offer. In all cases the ‘random’ technique was used to determine which goods should be requested in a bundle. Each case is labeled with its corresponding CATS legacy suite number; very similar distributions are given similar numbers and identical distributions are given the same number.

- [L1] *Sandholm*: Uniform, fixed random with *low\_fixed* = 0, *hi\_fixed* = 1
- [L1a] *Andersson et al.*: Uniform, fixed random with *low\_fixed* = 0, *hi\_fixed* = 1000
- [L2] *Sandholm*: Uniform, linearly random with *low\_linearly* = 0, *hi\_linearly* = 1
- [L2a] *Andersson et al.*: Uniform, linearly random with *low\_linearly* = 500, *hi\_linearly* = 1500
- [L3] *Sandholm*: Constant with *constant\_goods* = 3, fixed random with *low\_fixed* = 0, *hi\_fixed* = 1
- [L3] *deVries and Vohra*: Constant with *constant\_goods* = 3, fixed random with *low\_fixed* = 0, *hi\_fixed* = 1
- [L4] *Sandholm*: Decay with  $\alpha = 0.55$ , linearly random with *low\_linearly* = 0, *hi\_linearly* = 1
- [L4] *deVries and Vohra*: Decay with  $\alpha = 0.55$ , linearly random with *low\_linearly* = 0, *hi\_linearly* = 1
- [L4a] *Andersson et al.*: Decay with  $\alpha = 0.55$ , linearly random with *low\_linearly* = 1, *hi\_linearly* = 1000
- [L5] *Boutilier et al.*: Normal with  $\mu_{\text{goods}} = 4$  and  $\sigma_{\text{goods}} = 1$ , normal with  $\mu_{\text{price}} = 16$  and  $\sigma_{\text{price}} = 3$
- [L6] *Fujishima et al.*: Exponential with  $q = 5$ , linearly random with *low\_linearly* = 0.5, *hi\_linearly* = 1.5
- [L6a] *Andersson et al.*: Exponential with  $q = 5$ , linearly random with *low\_linearly* = 500, *hi\_linearly* = 1500
- [L7] *Fujishima et al.*: Binomial with  $p = 0.2$ , linearly random with *low\_linearly* = 0.5, *hi\_linearly* = 1.5
- [L7a] *Andersson et al.*: Binomial with  $p = 0.2$ , linearly random with *low\_linearly* = 500, *hi\_linearly* = 1500
- [L8] *deVries and Vohra*: Constant with *constant\_goods* = 3, quadratic

Parkes [17] used many of the test sets described above (particularly those described by Sandholm and Boutilier et

al.), but tested with fixed numbers of goods and bids rather than scaling these parameters.

### 5 CONCLUSION

In this paper we introduced CATS, a test suite for combinatorial auction optimization algorithms. The distributions in CATS represent a step beyond current CA testing techniques because they are economically motivated and model real-world problems. It is our hope that, with the help of others in the CA community, CATS will evolve into a universal test suite that will facilitate the development and evaluation of new CA optimization algorithms.

### 6 REFERENCES

- [1] A. Andersson, M. Tenhnen, , and F. Ygge. Integer programming for combinatorial auction winner determination. In *ICMAS-00*, 2000.
- [2] L. Ausubel, P. Crampton, R. McAfee, and J. McMillan. Synergies in wireless telephony: Evidence from the broadband PCS auctions. *Journal of Economics and Management Strategy*, 6(3):497–527, Fall 1997.
- [3] J. Banks, J. Ledyard, and D. Porter. Allocating uncertain and unresponsive resources: An experimental approach. *RAND Journal of Economics*, 20:1–23, 1989.
- [4] C. Boutilier, M. Goldszmidt, and B. Sabata. Sequential auctions for the allocation of resources with complementarities. Proceedings of IJCAI-99, 1999.
- [5] P. Brewer and C. Plott. A binary conflict ascending price (BICAP) mechanism for the decentralized allocation of the right to use railroad tracks. *International Journal of Industrial Organization*, 14:857–886, 1996.
- [6] M. Bykowsky, R. Cull, and J. Ledyard. Mutually destructive bidding: The FCC auction design problem. Technical Report Social Science Working Paper 916, California Institute of Technology, Pasadena, 1995.
- [7] C. DeMartini, A. Kwasnica, J. Ledyard, and D. Porter. A new and improved design for multi-object iterative auctions. Technical Report Social Science Working Paper 1054, California Institute of Technology, Pasadena, November 1998.
- [8] S. DeVries and R. Vohra. Combinatorial auctions: A survey. 2000.
- [9] W. Elmaghraby and S. Oren. The efficiency of multi-unit electricity auctions. *IAEE*, 20(4):89–116, 1999.
- [10] Y. Fujishima, K. Leyton-Brown, and Y. Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *IJCAI-99*, 1999.
- [11] D. Grether, R. Isaac, and C. Plott. *The Allocation of Scarce Resources: Experimental Economics and the Problem of Allocating Airport Slots*. Westview Press, Boulder, CO, 1989.
- [12] J. Ledyard, D. Porter, and A. Rangel. Experiments testing multiobject allocation mechanisms. *Journal of Economics & Management Strategy*, 6(3):639–675, Fall 1997.
- [13] J. Ledyard and K. Szakaly. Designing organizations

<sup>11</sup>Although in principle the problem of which goods to request could be answered in many ways, all legacy distributions of which we are aware use this technique.

<sup>12</sup>DeVries and Vohra [8] briefly describe a more general version of this price offer scheme, but do not describe how to set all the parameters (e.g., defining which goods are complementary); hence we do not include it here. Quadratic price offers may be particularly applicable to spectrum auctions; see [2].

- for trading pollution rights. *Journal of Economic Behavior and Organization*, 25:167–196, 1994.
- [14] D. Lehmann, L. O’Callaghan, and Y. Shoham. Truth revelation in rapid, approximately efficient combinatorial auctions. In *ACM Conference on Electronic Commerce*, 1999.
- [15] P. Milgrom. Putting auction theory to work: The simultaneous ascending auction. Technical Report 98-0002, Department of Economics, Stanford University, 1998.
- [16] N. Nisan. Bidding and allocation in combinatorial auctions. Working paper, 1999.
- [17] D. Parkes. ibundle: An efficient ascending price bundle auction. In *ACM Conference on Electronic Commerce*, 1999.
- [18] C. Plott and T. Cason. EPA’s new emissions trading mechanism: A laboratory evaluation. *Journal of Environmental Economics and Management*, 30:133–160, 1996.
- [19] D. Quan. Real estate auctions: A survey of theory and practice. *Journal of Real Estate Finance and Economics*, 9:23–49, 1994.
- [20] S. Rassenti, S. Reynolds, , and V. Smith. Cotenancy and competition in an experimental auction market for natural gas pipeline networks. *Economic Theory*, 4:41–65, 1994.
- [21] S. Rassenti, V. Smith, and R. Bulfin. A combinatorial auction mechanism for airport time slot allocation. *Bell Journal of Economics*, 13:402–417, 1982.
- [22] M. Rothkopf, A. Pekce, and R. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1998.
- [23] T. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. pages 256–262. Proceedings of AAAI-93, 1993.
- [24] T. Sandholm. An algorithm for optimal winner determination in combinatorial auctions. In *IJCAI-99*, 1999.
- [25] M. Tennenholtz. Some tractable combinatorial auctions. To appear in the proceedings of AAAI-2000, 2000.
- [26] M. Wellman, W. Walsh, P. Wurman, and J. MacKie-Mason. Auction protocols for decentralized scheduling. Proceedings of the 18th International Conference on Distributed Computing Systems, 1998.

## APPENDIX D

# R-MAX – A General Polynomial Time Algorithm for Near-Optimal Reinforcement Learning \*

**Ronen I. Brafman**  
*Computer Science Department  
Ben-Gurion University  
Beer-Sheva, Israel 84105*

BRAFMAN@CS.BGU.AC.IL

**Moshe Tennenholz**  
*Computer Science Department  
Stanford University  
Stanford, CA 94305*

MOSHE@ROBOTICS.STANFORD.EDU

### Abstract

R-MAX is an extremely simple model-based reinforcement learning algorithm which can attain near-optimal average reward in polynomial time. In R-MAX, the agent always maintains a complete, but possibly inaccurate model of its environment and acts based on the optimal policy derived from this model. The model is initialized in an optimistic fashion: all actions in all states return the maximal possible reward (hence the name). During execution, it is updated based on the agent’s observations. R-MAX improves upon several previous algorithms: (1) It is simpler and more general than Kearns and Singh’s  $E^3$  algorithm, covering zero-sum stochastic games. (2) It has a built-in mechanism for resolving the exploration vs. exploitation dilemma. (3) It formally justifies the “optimism under uncertainty” bias used in many RL algorithms. (4) It is simpler, more general, and more efficient than Brafman and Tennenholz’s LSG algorithm for learning in single controller stochastic games. (5) It generalizes the algorithm by Monderer and Tennenholz for learning in repeated games. (6) It is the only algorithm for learning in repeated games, to date, which is provably efficient, considerably improving and simplifying previous algorithms by Banos and by Megiddo.

### 1. Introduction

Reinforcement learning has attracted the attention of researchers in AI and related fields for quite some time. Many reinforcement learning algorithms exist and for some of them convergence rates are known. However, Kearns and Singh’s  $E^3$  algorithm (Kearns & Singh, 1998) was the first provably near-optimal polynomial time algorithm for learning in Markov decision processes (MDPs).  $E^3$  was extended later to handle single controller stochastic games (SCSGs) (Brafman & Tennenholz, 2000) as well as structured MDPs (Kearns & Koller, 1999). In  $E^3$  the agent learns by updating a model of its environment using statistics it collects. This learning process continues as long as it can be done relatively efficiently. Once this is no longer the case, the agent uses its learned model to compute an optimal

---

\*. This paper appeared as Technical Report 01-2001, Department of Computer Science, Ben-Gurion University. The second author permanent address is: Faculty of Industrial Engineering and Management, Technion-Israel Institute of Technology, Haifa 32000, Israel. The second author gratefully acknowledges the support of DARPA grant F30602-98-C-0214-P00005.

policy and follows it. The success of this approach rests on two important properties: the agent can determine online whether an efficient learning policy exists, and if such a policy does not exist, it is guaranteed that the optimal policy with respect to the learned model will be approximately optimal with respect to the real world.

The difficulty in generalizing  $E^3$  to adverserial contexts, i.e., to different classes of games, stems from the adversary's ability to influence the probability of reaching different states. In a game, the agent does not control its adversary's choices, nor can it predict them with any accuracy. Therefore, it has difficulty predicting the outcome of its actions and whether or not they will lead to new information. Consequently, it is unlikely that an agent can explicitly choose between an exploration and an exploitation policy. For this reason, the only extension of  $E^3$  to adverserial contexts used the restricted SCSG model in which the adversary influences the reward of a game only, and not its dynamics.

To overcome this problem, we suggest a different approach in which the agent never attempts to learn explicitly. Our agent always attempts to optimize its behavior, albeit with respect to a fictitious model in which optimal behavior often leads to learning. This model assumes that the reward the agent obtains in any situation it is not too familiar with, is the maximal possible reward –  $R_{max}$ . The optimal policy with respect to the agent's *fictitious* model has a very interesting and useful property with respect to the *real* model: it is always either optimal or it leads to efficient learning. The agent does *not* know whether it is optimizing or learning efficiently, but it always does one or the other. Thus, the agent will always either exploit or explore efficiently, without knowing ahead of time which of the two will occur. Since there is only a polynomial number of parameters to learn, as long as learning is done efficiently we can ensure that the agent spends a polynomial number of steps exploring, and the rest of the time will be spent exploiting. Thus, the resulting algorithm may be said to use an *implicit* explore or exploit approach, as opposed to Kearns and Singh's *explicit* explore or exploit approach.

This learning algorithm, which we call R-MAX, is very simple to understand and to implement. The algorithm converges in polynomial-time to a near-optimal solution. Moreover, R-MAX is described in the context of zero-sum stochastic game, a model that is more general than Markov Decision Processes. As a consequence, R-MAX is more general and more efficient than a number of previous results. It generalizes the results of Kearns and Singh (1998) to adverserial contexts and to situations where the agent considers a stochastic model of the environment inappropriate, opting for a non-deterministic model instead. R-MAX can handle more classes of stochastic games than the LSG algorithm (Brafman & Tennenholtz, 2000). In addition, it attains a higher expected average reward than LSG. R-MAX also improves upon previous algorithms for learning in repeated games (Aumann & Maschler, 1995), such as Megiddo's (Megiddo, 1980) and Banos (Banos, 1968). It is the only polynomial time algorithm for this class of games that we know of, and it is much simpler, too. Finally, R-MAX generalizes the results of Monderer and Tennenholtz (Monderer & Tennenholtz, 1997) to handle the general probabilistic maximin (safety level) decision criterion.

The approach taken by R-MAX is not new. It has been referred to as *the optimism in the face of uncertainty* heuristic, and was considered an ad-hoc, though useful, approach (e.g., see Section 2.2.1 in (Kaelbling, Littman, & Moore, 1996), where it appears under the heading "Ad-Hoc Techniques" and Section 2.7 in (Sutton & Barto, 1998) where this approach is

called *optimistic initial values* and is referred to as a “simple trick that can be quite effective on stationary problems”). This optimistic bias has been used in a number of well-known reinforcement learning algorithms, e.g. Kaelbling’s interval exploration method (Kaelbling, 1993), the exploration bonus in Dyna (Sutton, 1990), the curiosity-driven exploration of (Schmidhuber, 1991), and the exploration mechanism in prioritized sweeping (Moore & Atkeson, 1993). More recently, Tadepalli and Ok (Tadepalli & Ok, 1998) presented a reinforcement learning algorithm that works in the context of the undiscounted average-reward model used in this paper. In particular, one variant of their algorithm, called AH-learning, is very similar to R-MAX. However, as we noted above, none of this work provides theoretical justification for this very natural bias. Thus, an additional contribution of this paper is a formal justification for the *optimism under uncertainty* bias.

The paper is organized as follows: in Section 2 we define the learning problem more precisely and the relevant parameters. In Section 3 we describe the R-MAX algorithm. In Section 4 we prove that it yields near-optimal reward in polynomial time. We conclude in Section 5.

## 2. Preliminaries

We present R-MAX in the context of a model that is called a *stochastic game*. This model is more general than a Markov decision process because it does not necessarily assume that the environment acts stochastically (although it can). In what follows we define the basic model, describe the set of assumptions under which our algorithm operates, and define the parameters influencing its running time.

### 2.1 Stochastic Games

A game is a model of multi-agent interaction. In a game, we have a set of players, each of whom chooses some action to perform from a given set of actions. As a result of the players’ combined choices, some outcome is obtained which is described numerically in the form of a payoff vector, i.e., a vector of values, one for each of the players. We concentrate on two-player, fixed-sum games (i.e., games in which the sum of values in the payoff vector is constant). We refer to the player under our control as *the agent*, whereas the other player will be called *the adversary*.

A common description of a game is as a matrix. This is called a game in *strategic form*. The rows of the matrix correspond to the agent’s actions and the columns correspond to the adversary’s actions. The entry in row  $i$  and column  $j$  in the game matrix contains the rewards obtained by the agent and the adversary if the agent plays his  $i^{th}$  action and the adversary plays his  $j^{th}$  action. We make the simplifying assumption that the size of the action set of both the agent and the adversary is identical. However, an extension to sets of different sizes is trivial.

In a *stochastic game* (SG) the players play a (possibly infinite) sequence of standard games from some given set of games. After playing each game, the players receive the appropriate payoff, as dictated by that game’s matrix, and move to a new game. The identity of this new game depends, stochastically, on the previous game and on the players’ actions in that previous game. Formally:

**Definition 1** A fixed-sum, two player, stochastic-game [SG]  $M$  on states  $S = \{1, \dots, N\}$ , and actions  $A = \{a_1, \dots, a_k\}$ , consists of:

- **Stage Games:** each state  $s \in S$  is associated with a two-player, fixed-sum game in strategic form, where the action set of each player is  $A$ . We use  $R^i$  to denote the reward matrix associated with stage-game  $i$ .
- **Probabilistic Transition Function:**  $P_M(s, t, a, a')$  is the probability of a transition from state  $s$  to state  $t$  given that the first player (the agent) plays  $a$  and the second player (the adversary) plays  $a'$ .

An SG is similar to an MDP. In both models, actions lead to transitions between states of the world. The main difference is that in an MDP the transition depends on the action of a single agent whereas in an SG the transition depends on a joint-action of the agent and the adversary. In addition, in an SG, the reward obtained by the agent for performing an action depends on its action *and* the action of the adversary. To model this, we associate a game with every state. Therefore, we shall use the terms *state* and *game* interchangeably.

Stochastic games are useful not only in multi-agent contexts. They can be used instead of MDPs when we do not wish to model the environment (or certain aspects of it) stochastically. In that case, we can view the environment as an agent that can choose among different alternatives, without assuming that its choice is based on some probability distribution. This leads to behavior maximizing the worst-case scenario. In addition, the adversaries that the agent meets in each of the stage-games could be different entities.

R-MAX is formulated as an algorithm for learning in Stochastic Games. However, it is immediately applicable to fixed-sum repeated games and to MDPs because both of these models are degenerate forms of SGs. A repeated game is an SG with a single state and an MDP is an SG in which the adversary has a single action at each state.

For ease of exposition we normalize both players' payoffs in each stage game to be non-negative reals between 0 and some constant  $R_{max}$ . We also take the number of actions to be constant. The set of possible histories of length  $t$  is  $(S \times A^2)^t \times S$ , and the set of possible histories,  $H$ , is the union of the sets of possible histories for all  $t \geq 0$ , where the set of possible histories of length 0 is  $S$ .

Given an SG, a policy for the agent is a mapping from  $H$  to the set of possible probability distributions over  $A$ . Hence, a policy determines the probability of choosing each particular action for each possible history.

We define the *value* of a policy using the *average expected reward criterion* as follows: Given an SG  $M$  and a natural number  $T$ , we denote the expected  $T$ -step undiscounted average reward of a policy  $\pi$  when the adversary follows a policy  $\rho$ , and where both  $\pi$  and  $\rho$  are executed starting from a state  $s \in S$ , by  $U_M(s, \pi, \rho, T)$  (we omit subscripts denoting the SG when this causes no confusion). Let  $U_M(s, \pi, T) = \min_{\rho \text{ is a policy}} U_M(s, \pi, \rho, T)$  denote the value that a policy  $\pi$  can guarantee in  $T$  steps starting from  $s$ . We define  $U_M(s, \pi) = \liminf_{T \rightarrow \infty} U_M(s, \pi, T)$ . Finally, we define  $U_M(\pi) = \min_{s \in S} U_M(s, \pi)$ .<sup>1</sup>

---

1. We discuss this choice below.

## 2.2 Assumptions, Complexity and Optimality

We make two central assumptions: First, we assume that the agent always recognizes the identity of the state (or stage-game) it reached (but not its associated payoffs and transition probabilities) and that after playing a game, it knows what actions were taken by its adversary and what payoffs were obtained. Second, we assume that the maximal possible reward  $R_{max}$  is known ahead of time. This latter assumption can be removed.<sup>2</sup>

Next, we wish to discuss the central parameter in the analysis of the complexity of R-MAX – the *mixing time*, first identified by Kearns and Singh (1998). Kearns and Singh argue that it is unreasonable to refer to the efficiency of learning algorithms without referring to the efficiency of convergence to a desired value. They defined the  $\epsilon$ -return mixing time of a policy  $\pi$  to be the smallest value of  $T$  after which  $\pi$  guarantees an expected payoff of at least  $U(\pi) - \epsilon$ . In our case, we have to take into account the existence of an adversary. Therefore, we adjust this definition slightly as follows: a policy  $\pi$  belongs to the set  $\Pi(\epsilon, T)$  of policies whose  $\epsilon$ -return mixing time is at most  $T$ , if for any starting state  $s$  and for any adversary behavior  $\rho$ , we have that  $U(s, \pi, \rho, T) > U(\pi) - \epsilon$ .

That is, if a policy  $pi \in \Pi(\epsilon, T)$  then no matter what the initial state is and what the adversary does, the policy  $\pi$  will yield in any  $t \geq T$  steps an expected average reward that is  $\epsilon$  close to its value. The  $\epsilon$ -return mixing time of a policy  $\pi$  is the smallest  $T$  for which  $pi \in \Pi(\epsilon, T)$ . Notice that this means that an agent with perfect information about the nature of the games and the transition function will require at least  $T$  steps, on the average, to obtain an optimal value using an optimal policy  $\pi$  whose  $\epsilon$ -return mixing time is  $T$ . Clearly, one cannot expect an agent lacking this information to perform better.

We denote by  $Opt(\Pi(\epsilon, T))$  the optimal expected  $T$ -step undiscounted average return from among the policies in  $\Pi(\epsilon, T)$ . When looking for an optimal policy (with respect to policies that mix at time  $T$ , for a given  $\epsilon > 0$ ), we will be interested in approaching this value in time polynomial in  $T$ , in  $1/\epsilon$ , in  $1/\delta$  (where  $\epsilon$  and  $\delta$  are the desired error bounds), and in the size of the description of the game.

The reader may have noticed that we defined  $U_M(\pi)$  as  $\min_{s \in S} U_M(s, \pi)$ . It may appear that this choice makes the learning task too easy. For instance, one may ask why shouldn't we try to attain the maximal value over all possible states, or at least the value of our initial state? We claim that the above is the only reasonable choice, and that it leads to results that are as strong as previous algorithms.

To understand this point, consider the following situation: we start learning at some state  $s$  in which the optimal action is  $a$ . If we do not execute the action  $a$  in  $s$ , we reach some state  $s'$  that has a very low value. A learning algorithm without any prior knowledge cannot be expected to immediately guess that  $a$  should be done in  $s$ . In fact, without such prior knowledge, it cannot conclude that  $a$  is a good action unless it tries the other actions in  $s$  and compares their outcome to that of  $a$ . Thus, one can expect an agent to learn a near-optimal policy only if the agent can visit state  $s$  sufficiently many times to learn about the consequences of different options in  $s$ . In a finite SG, there will be some set of states that we can sample sufficiently many times, and it is for such states that we can learn to behave.

---

2. We would need to run the algorithm repeatedly for increasing values of  $R_{max}$ . The resulting algorithm remains polynomials in the relevant parameters.

In fact, it probably makes sense to restrict our attention to a subset of the states such that from each state in this set it is not too hard to get to any other state. In the context of MDPs, Kearns and Singh refer to this as the *ergodicity* assumption. In the context of SGs, Hoffman and Karp (1966) refer to this as the *irreducibility* assumption. An SG is said to be *irreducible* if the Markov-chain obtained by fixing any two (pure) stationary strategies for each of the players is irreducible (i.e., each state is reachable from each other state). In the special case of an MDP, irreducibility is precisely the ergodicity property used by Kearns and Singh in their analysis of  $E^3$ .

Irreducible SGs have a number of nice properties, as shown by (Hoffman & Karp, 1966). First, the maximal long-term average reward is independent of the starting state, implying that  $\max_{\pi} \min_{s \in S} U_M(s, \pi) = \max_{\pi} \max_{s \in S} U_M(s, \pi)$ . Second, this optimal value can be obtained by a stationary policy (i.e., one that depends on the current stage-game only). Thus, although we are not restricting ourselves to irreducible games, we believe that our results are primarily interesting in this class of games.

### 3. The R-MAX algorithm

Recall that we consider a stochastic game  $M$  consisting of a set  $S = \{G_1, \dots, G_N\}$  of stage-games in each of which both the agent and the adversary have a set  $A = \{a_1, \dots, a_k\}$  of possible actions. We associate a reward matrix  $R^i$  with each game, and use  $R_{m,l}^i$  to denote a pair consisting of the reward obtained by the agent and the adversary after playing actions  $a_m$  and  $a_l$  in game  $G_i$ , respectively. In addition, we have a probabilistic transition function,  $P_M$ , such that  $P_M(s, t, a, a')$  is the probability of making a transition from  $G_s$  to  $G_t$  given that the agent played  $a$  and the adversary played  $a'$ . It is convenient to think of  $P_M(i, \cdot, a, a')$  as a function associated with the entry  $(a, a')$  in the stage-game  $G_i$ . This way, all model parameters, both rewards and transitions, are associated with joint actions of a particular game. Let  $\epsilon > 0$ . For ease of exposition, we assume throughout most of the analysis that the  $\epsilon$ -return mixing time of the optimal policy,  $T$ , is known. Later, we show how this assumption can be relaxed.

The R-MAX algorithm is defined as follows:

**Initialize:** Construct the following model  $M'$  consisting of  $N+1$  stage-games,  $\{G_0, G_1, \dots, G_N\}$ , and  $k$  actions,  $\{a_1, \dots, a_k\}$ . Here,  $G_1, \dots, G_N$  correspond to the real games,  $\{a_1, \dots, a_k\}$  correspond to the real actions, and  $G_0$  is an additional fictitious game. Initialize all game matrices to have  $(R_{max}, 0)$  in all entries.<sup>3</sup> Initialize  $P_M(G_i, G_0, a, a') = 1$  for all  $i = 0, \dots, N$  and for all actions  $a, a'$ .

In addition, maintain the following information for each entry in each game  $G_1, \dots, G_N$ : (1) a boolean value *known/unknown*, initialized to *unknown*; (2) the states reached by playing the joint action corresponding to this entry (and how many times); (3) the reward obtained (by both players) when playing the joint action corresponding to this entry. Items 2 and 3 are initially empty.

**Repeat:**

---

3. The value 0 given to the adversary does not play an important role here.

**Compute and Act:** Compute an optimal  $T$ -step policy for the current state, and execute it for  $T$ -steps or until a new entry becomes known.

**Observe and update:** Following each joint action do as follows: Let  $a$  be the action you performed in  $G_i$  and let  $a'$  be the adversary's action.

- If the joint action  $(a, a')$  is performed for the first time in  $G_i$ , update the reward associated with  $(a, a')$  in  $G_i$ , as observed.
- Update the set of states reached by playing  $(a, a')$  in  $G_i$ .
- If at this point your record of states reached from this entry contains  $K_1 = \max((\lceil \frac{4NTR_{\max}}{\epsilon} \rceil^3), \lceil -6\ln^3(\frac{\delta}{6Nk^2}) \rceil) + 1$  elements, mark this entry as *known*, and update the transition probabilities for this entry according to the observed frequencies.

As can be seen, R-MAX is quite simple. It starts with an initial estimate for the model parameters that assumes all states and all joint actions yield maximal reward and lead with probability 1 to the fictitious stage-game  $G_0$ . Based on the current model, an optimal policy is computed and followed. Following each joint action the agent arrives at a new stage-game, and this transition is recorded in the appropriate place. Once we have enough information about where some joint action leads to from some stage-game, we update the entries associated with this stage-game and this joint action in our model. After each model update, we recompute an optimal policy and repeat the above steps.

#### 4. Optimality and Convergence

In this section we provide the tools that ultimately lead to the proof of the following theorem:

**Theorem 1** *Let  $M$  be an SG with  $N$  states and  $k$  actions. Let  $0 < \delta < 1$ , and  $\epsilon > 0$  be constants. Denote the policies for  $M$  whose  $\epsilon$ -return mixing time is  $T$  by  $\Pi_M(\epsilon, T)$ , and denote the optimal expected return achievable by such policies by  $Opt(\Pi_M(\epsilon, T))$ . Then, with probability of no less than  $1 - \delta$  the R-MAX algorithm will attain an expected return of  $Opt_M(\Pi(\epsilon, T)) - 2\epsilon$  within a number of steps polynomial in  $N, k, T, \frac{1}{\epsilon}$ , and  $\frac{1}{\delta}$ .*

In the main lemma required for proving this theorem we show the following: if the agent follows a policy that is optimal with respect to the model it maintains for  $T$  steps, it will either attain near-optimal average reward, as desired, or it will update its statistics for one of the unknown slots with sufficiently high probability. This can be called the *implicit explore or exploit* property of R-MAX: The agent does not know ahead of time whether it is exploring or exploiting – this depends in a large part on the adversary's behavior which it cannot control or predict. However, it knows that it does one or the other, no matter what the adversary does. Using this result we can proceed as follows: As we will show, the number of samples required to mark a slot as known is polynomial in the problem parameters, and so is the total number of entries. Therefore, the number of  $T$ -step iterations in which non-optimal reward is obtained is bounded by some polynomial function of the input parameters, say  $T'$ . This implies that by performing  $T$ -step iterations  $D = T'R_{\max}/\theta$  times, we get that the loss obtained by non-optimal execution (where exploration is performed), is bounded by  $\theta$ , for any  $0 < \theta < 1$ .

Before proving our main lemma we state and prove an extension of Kearns and Singh's Simulation Lemma (Kearns & Singh, 1998) to the context of SGs with a slightly improved bound.

**Definition 2** Let  $M$  and  $\bar{M}$  be SGs over the same state and action spaces. We say that  $\bar{M}$  is an  $\alpha$ -approximation of  $M$  if for every state  $s$  we have:

1. If  $P_M(s, t, a, a')$  and  $P_{\bar{M}}(s, t, a, a')$  are the probabilities of transition from state  $s$  to state  $t$  given that the joint action carried out by the agent and the adversary is  $(a, a')$ , in  $M$  and  $\bar{M}$  respectively, then,  $P_M(s, t, a, a') - \alpha \leq P_{\bar{M}}(s, t, a, a') \leq P_M(s, t, a, a') + \alpha$
2. For every state  $s$ , the same stage-game is associated with  $s$  in  $\bar{M}$  and in  $M$ .

**Lemma 1** Let  $M$  and  $\bar{M}$  be SGs over  $N$  states, where  $\bar{M}$  is an  $\frac{\epsilon}{NTR_{max}}$ -approximation of  $M$ , then for every state  $s$ , agent policy  $\pi$ , and adversary policy  $\rho$ , we have that

$$|U_{\bar{M}}(s, \pi, \rho, T) - U_M(s, \pi, \rho, T)| \leq \epsilon.$$

**Proof:** When we fix both players' policies we get, both in MDPs and in general SGs, a probability distribution over  $T$ -step paths in the state space. This is not a Markov process because the player's policies can be non-stationary. However, the transition probabilities at each point depend on the current state and the actions taken and the probability of each path is a product of the probability of each of the transitions. This is true whether the policies are pure or mixed.

We need to prove that:

$$\sum_p |\Pr_M(p)U_M(p) - \Pr_{\bar{M}}(p)U_{\bar{M}}(p)| \leq \epsilon$$

where  $p$  is a  $T$ -step path starting at  $s$ ,  $\Pr_M(p)$  (respectively,  $\Pr_{\bar{M}}(p)$ ) is its probability in the random process induced by  $M$  (resp. by  $\bar{M}$ ),  $\pi$ , and  $\rho$ , and  $U_M(p)$ , ( $U_{\bar{M}}(p)$ ) is the average payoff along this path. Because the average payoff is bound by  $R_{max}$  we have:

$$\sum_p |\Pr_M(p)U_M(p) - \Pr_{\bar{M}}(p)U_{\bar{M}}(p)| \leq \sum_p |\Pr_M(p) - \Pr_{\bar{M}}(p)|R_{max}.$$

To conclude our proof, it is sufficient to show that

$$\sum_p |\Pr_M(p) - \Pr_{\bar{M}}(p)| \leq \epsilon/R_{max}$$

Let  $h_i$  define the following random processes: start at state  $s$  and follow policies  $\rho$  and  $\tau$ ; for the first  $i$  steps, the transition probabilities are identical to the process defined above on  $\bar{M}$ , and for the rest of the steps its transition probabilities are identical to  $M$ . Clearly, when we come to assess the probabilities of  $T$ -step path, we have that  $h_0$  is identical to the original process on  $M$ , whereas  $h_T$  is identical to original process on  $\bar{M}$ . The triangle inequality implies that

$$\sum_p |\Pr_M(p) - \Pr_{\bar{M}}(p)| = \sum_p |\Pr_M(p) - \Pr_{h_T}(p)| \leq \sum_{i=0}^{T-1} \sum_p |\Pr_{h_i}(p) - \Pr_{h_{i+1}}(p)|$$

If we show that for any  $0 \leq i < T$  we have that  $\sum_p |\Pr_{h_i}(p) - \Pr_{h_{i+1}}(p)| \leq \epsilon/TR_{max}$ , it will follow that  $\sum_p |\Pr_M(p) - \Pr_{\bar{M}}(p)| \leq \epsilon/R_{max}$ , which is precisely what we need to show.

We are left with the burden of proving that  $\sum_p |\Pr_{h_i}(p) - \Pr_{h_{i+1}}(p)| \leq \epsilon/TR_{max}$ . We can sum over all path  $p$  as follows: first we sum over the  $N$  possible states that can be reached in  $i$  steps. Then we sum over all possible path prefixes that reach each such state. Next, we sum over all possible states reached after step  $i+1$ , and finally over all possible suffixes that start from each such state. Now, we note that the probability of each particular path  $p$  is the product of the probability of its particular prefix, the probability of a transition from  $x_i$  to  $x_{i+1}$ , and the probability of the suffix. We will use  $x_i$  to denote the state reached after  $i$  steps,  $x_{i+1}$  to denote the state reached after  $i+1$  steps,  $pre(x_i)$  to denote the  $i$ -step prefixes reaching  $x_i$ , and  $suf(x_j)$  to denote the suffixes starting at  $x_j$ . Thus,

$$\begin{aligned} \sum_p |\Pr_{h_i}(p) - \Pr_{h_{i+1}}(p)| &= \sum_{x_i} \sum_{pre(x_i)} \sum_{x_{i+1}} \sum_{suf(x_{i+1})} |(\Pr_{h_i}(pre(x_i)) \Pr_{h_i}(x_i \rightarrow x_{i+1}) \Pr_{h_i}(suf(x_{i+1}))) - \\ &\quad (\Pr_{h_{i+1}}(pre(x_i)) \Pr_{h_{i+1}}(x_i \rightarrow x_{i+1}) \Pr_{h_{i+1}}(suf(x_{i+1})))| \end{aligned}$$

However, the prefix and suffix probabilities are identical in  $h_i$  and  $h_{i+1}$ . Thus, this sum is equal to

$$\begin{aligned} \sum_{x_i} \sum_{pre(x_i)} \sum_{x_{i+1}} \sum_{suf(x_{i+1})} \Pr_{h_i}(pre(x_i)) \Pr_{h_i}(suf(x_{i+1})) |\Pr_{h_i}(x_i \rightarrow x_{i+1}) - \Pr_{h_{i+1}}(x_i \rightarrow x_{i+1})| = \\ \sum_{x_i} \sum_{pre(x_i)} \Pr_{h_i}(pre(x_i)) \sum_{x_{i+1}} \sum_{suf(x_{i+1})} \Pr_{h_i}(suf(x_{i+1})) |\Pr_{h_i}(x_i \rightarrow x_{i+1}) - \Pr_{h_{i+1}}(x_i \rightarrow x_{i+1})| \leq \\ [\sum_{x_i} \sum_{pre(x_i)} \Pr_{h_i}(pre(x_i))] [\sum_{x_{i+1}} \sum_{suf(x_{i+1})} \Pr_{h_i}(suf(x_{i+1}))] \epsilon/NTR_{max} \end{aligned}$$

This last expression is a product of two independent terms. The first term is the sum over all possible  $i$ -step prefixes (i.e., overall all prefixes starting in the given  $x_0$  and ending in  $x_i$ , for any  $x_i$ ). Hence, it is equal to 1. The second term is a sum over all suffixes starting at  $x_{i+1}$ , for any value of  $x_{i+1}$ . For any given value of  $x_{i+1}$  the probability of any suffix starting at this value is 1. Summing over all possible values of  $x_{i+1}$ , we get a value of  $N$ .

Thus,

$$\sum_p |\Pr_{h_i}(p) - \Pr_{h_{i+1}}(p)| \leq 1 \cdot \epsilon/NTR_{max} \cdot N$$

This concludes our proof.  $\blacksquare$

Next, we define the notion of an *induced SG*. The definition is similar to the definition of an induced MDP given in (Kearns & Singh, 1998) except for the use of R-MAX. The induced SG is the model used by the agent to determine its policy.

**Definition 3** Let  $M$  be an SG. Let  $L$  be the set of entries  $(G_i, a, a')$  marked unknown. That is, if  $(G_i, a, a') \in L$  then the entry corresponding to the joint action  $(a, a')$  in the stage-game  $G_i$  is marked as unknown. Define  $M_L$  to be the following SG:  $M_L$  is identical to  $M$ , except that  $M_L$  contains an additional state  $G_0$ . Transitions and rewards associated

with all entries in  $M_L$  which are not in  $L$  are identical to those in  $M$ . For any entry in  $L$  or in  $G_0$ , the transitions are with probability 1 to  $G_0$ , and the reward is  $R_{max}$  for the agent and 0 for the adversary.

Given an SG  $M$  with a set  $L$  of unknown states,  $R^{M_L}$ -max denotes the optimal policy for the induced SG  $M_L$ . When  $M_L$  is clear from the context we will simply use the term *R-max policy* instead of  $R^{M_L}$ -max policy.

We now state and prove the implicit explore or exploit lemma:

**Lemma 2** *Let  $M$  be an SG, let  $L$  and  $M_L$  be as above. Let  $\rho$  be an arbitrary policy for the adversary, let  $s$  be some state, and let  $0 < \alpha < 1$ . Then either (1)  $|Opt(\Pi_M(\epsilon, T)) - V_{R\text{-max}}| < \alpha$ , where  $V_{R\text{-max}}$  is the expected  $T$ -step average reward for the  $R^{M_L}$ -max policy on  $M$ ; or (2) An unknown entry will be played in the course of running R-max on  $M$  for  $T$  steps with a probability of at least  $\frac{\alpha}{R_{max}}$ .*

In practice, we cannot determine  $\rho$ , the adversary's policy, ahead of time. Thus, we do not know whether R-max will attain near-optimal reward or whether it will reach an unknown entry with sufficient probability. The crucial point is that it will do one or the other, no matter what the adversary does.

**Proof:** First, notice that the value of R-MAX in  $M_L$  will be no less than the value of the optimal policy in  $M$ . This follows from the fact that the reward for the agent is at least as large as in  $M$ , and that the R-max policy is optimal with respect to  $M_L$ .

In order to prove the claim, we will show that the difference between the reward obtained by the agent in  $M$  and in  $M_L$  when R-MAX is played is smaller than the exploration probability times  $R_{max}$ . This will imply that if the exploration probability is small, then R-MAX will attain near-optimal payoff. Conversely, if near-optimal payoff is not attained, the exploration probability will be sufficiently large.

For any policy, we may write:

$$U_M(s, \pi, \rho, T) = \sum_p Pr_M^{\pi, s}[p] U_M(p) = \sum_q Pr_M^{\pi, s}[q] U_M(q) + \sum_r Pr_M^{\pi, s}[r] U_M(r)$$

where the sums are over, respectively, all  $T$ -paths  $p$  in  $M$ , all  $T$ -paths  $q$  in  $M$  such that every entry visited in  $q$  is not in  $L$ , and all  $T$ -path  $r$  in  $M$  in which at least one entry visited is in  $L$ . Hence:

$$\begin{aligned} |U_M(s, \text{R-max}, \rho, T) - U_{M_L}(s, \text{R-max}, \rho, T)| &= |\sum_p Pr_M^{\text{R-max}, \rho, s}[p] U_M(p) - \sum_p Pr_{M_L}^{\text{R-max}, \rho, s}[p] U_{M_L}(p)| \\ &= |\sum_q Pr_M^{\text{R-max}, \rho, s}[q] U_M(q) + \sum_r Pr_M^{\text{R-max}, \rho, s}[r] U_M(r) - \sum_q Pr_{M_L}^{\text{R-max}, \rho, s}[q] U_{M_L}(q) + \sum_r Pr_{M_L}^{\text{R-max}, \rho, s}[r] U_{M_L}(r)| \\ &\leq |\sum_q Pr_M^{\text{R-max}, \rho, s}[q] U_M(q) - \sum_q Pr_{M_L}^{\text{R-max}, \rho, s}[q] U_{M_L}(q)| + \\ &\quad |\sum_r Pr_M^{\text{R-max}, \rho, s}[r] U_M(r) - \sum_r Pr_{M_L}^{\text{R-max}, \rho, s}[r] U_{M_L}(r)| \end{aligned}$$

The first difference:

$$|\sum_q \Pr_M^{\text{R-max}, \rho, s}[q] U_M(q) - \sum_q \Pr_{M_L}^{\text{R-max}, \rho, s}[q] U_{M_L}(q)|$$

must be 0. This follows from the fact that in  $M$  and in  $M_L$ , the rewards obtained in a path which do not visit an unknown entry are identical. The probability of each such path is identical as well.

Hence, we have:

$$\begin{aligned} |U_M(s, \text{R-max}, \rho, T) - U_{M_L}(s, \text{R-max}, \rho, T)| &\leq |\sum_r \Pr_M^{\text{R-max}, \rho, s}[r] U_M(r) - \sum_r \Pr_{M_L}^{\text{R-max}, \rho, s}[r] U_{M_L}(r)| \\ &\leq \sum_r \Pr_M^{\text{R-max}, \rho, s}[r] R_{\max} \end{aligned}$$

This last inequality stems from the fact that the average reward in any path is no greater than  $R_{\max}$  and no smaller than 0 and the fact that we can appropriately associate different paths within these models with equal probabilities.

The last term is the probability of reaching an unknown entry multiplied by  $R_{\max}$ . If this probability is less than  $\frac{\alpha}{R_{\max}}$  then

$$|U_M(s, \text{R-max}, \rho, T) - U_{M_L}(s, \text{R-max}, \rho, T)| \leq \alpha$$

Denote by  $\pi^*$  an optimal  $T$ -step policy, and let  $U_M(s, \pi, T)$  be its value. (Note that this value is independent of the adversary strategy  $\rho$ , as  $\pi^*$  guarantees at least this value for every adversary behavior.) If  $U_M(s, \text{R-max}, \rho, T) \geq U_M(s, \pi, T)$  we are done. Suppose that  $U_M(s, \text{R-max}, \rho, T) < U_M(s, \pi, T)$ . We know that  $U_{M_L}(s, \text{R-max}, \rho, T) \geq U_M(s, \pi, T)$  is no lesser than the optimal  $T$  step average reward for  $M$ . Therefore, we have that

$$\begin{aligned} |U_M(s, \pi, T) - U_M(s, \text{R-max}, \rho, T)| &= U_M(s, \pi, T) - U_M(s, \text{R-max}, \rho, T) \\ &\leq U_{M_L}(s, \text{R-max}, \rho, T) - U_M(s, \text{R-max}, \rho, T) \leq \alpha \end{aligned}$$

■

We are now ready to prove Theorem 1. First, we wish to show that the expected average reward is as stated. We must consider three models:  $M$ , the real model,  $M'_L$  the actual model used, and  $M'$  where  $M'$  is an  $\epsilon/2NTR_{\max}$ -approximation of  $M$  such that the SG induced by  $M'$  and  $L$  is  $M'_L$ . At each  $T$ -step iteration of our algorithm we can apply the Implicit Explore or Exploit Lemma to  $M'$  and  $M'_L$  for the set  $L$  applicable at that stage. Hence, at each step either the current R-max policy leads to an average reward that is  $\epsilon/2$  close to optimal with respect to the adversary's behavior and the model  $M'$  or it leads to an efficient learning policy with respect to the same model. However, because  $M'$  is an  $\epsilon/2$ -approximation of  $M$ , the simulation lemma guarantees that the policy generated is either  $\epsilon$  close to optimal or explores efficiently. We know that the number of  $T$ -step phases in which we are exploring can be bounded polynomially. This follows from the fact that we have a polynomial number of parameters to learn (in  $N$  and  $k$ ) and that the probability that we obtain a new, useful statistic is polynomial in  $\epsilon, T$  and  $N$ . Thus, if we choose a

large enough (but still polynomial) number of  $T$ -step phases, we shall guarantee that our average reward is as close to optimal as we wish.

The above analysis was done assuming we actually obtain the expected value of each random variable. This cannot be guaranteed with probability 1. Yet, we can ensure that the probability that the algorithm fails to attain the expected value of certain parameters be small enough by sampling it a larger (though still polynomial) number of time. This is based on the well-known Chernoff bound. Using this technique one can show that when the variance of some random variable is bounded, we can ensure that we get near its average with probability  $1 - \delta$  by using a sufficiently large sample that is polynomial in  $1/\delta$ .

In our algorithm, there are three reasons why the algorithm could fail to provided the agent with near optimal return in polynomial time.

1. First, we have to guarantee that our estimates of the transition probabilities for every slot are sufficiently accurate. Recall that to ensure a loss of no more than  $\epsilon/2$  our estimates must be within  $\frac{\epsilon}{2NTR_{max}}$  of the real probabilities.

Consider a set of trials, where the joint action  $(a, a')$  is performed in state  $s$ . Consider the probability of moving from state  $s$  to state  $t$  given the joint-action  $(a, a')$  in a given trial, and denote it by  $p$ . Notice that there are  $Nk^2$  such probabilities (one for each game and pair of agent-adversary actions). Therefore, we would like to show that the probability of failure in estimating  $p$  is less than  $\frac{\delta}{3Nk^2}$ . Let  $X_i$  be an indicator random variable, that is 1 iff we moved to state  $t$  when we were in state  $s$  and selected an action  $a$  in trial  $i$ . Let  $Z_i = X_i - p$ . Then  $E(Z_i) = 0$ , and  $|Z_i| \leq 1$ .

Then, Chernoff bound implies that (for any  $K_1$ )  $Prob(\sum_{i=1}^{K_1} Z_i > K_1^{\frac{2}{3}}) < e^{-\frac{K_1^{\frac{1}{3}}}{2}}$ . This implies that  $Prob(\frac{\sum_{i=1}^{K_1} X_i}{K_1} - p > K_1^{-\frac{1}{3}}) < e^{-\frac{K_1^{\frac{1}{3}}}{2}}$ . Similarly, we can define  $Z'_i = p - X_i$ , and get by Chernoff bound that  $Prob(\sum_{i=1}^{K_1} Z'_i > K_1^{\frac{2}{3}}) < e^{-\frac{K_1^{\frac{1}{3}}}{2}}$ . This implies that  $Prob(p - \frac{\sum_{i=1}^{K_1} X_i}{K_1} > K_1^{-\frac{1}{3}}) < e^{-\frac{K_1^{\frac{1}{3}}}{2}}$ . Hence, we get that  $Prob(|\frac{\sum_{i=1}^{K_1} X_i}{K_1} - p| > K_1^{-\frac{1}{3}}) < 2e^{-\frac{K_1^{\frac{1}{3}}}{2}}$ .

We now choose  $K_1$  such that  $K_1^{-\frac{1}{3}} < \frac{\epsilon}{2NTR_{max}}$ , and  $2e^{-\frac{K_1^{\frac{1}{3}}}{2}} < \frac{\delta}{3Nk^2}$ . This is obtained by taking  $K_1 = \max((\frac{4NTR_{max}}{\epsilon})^3, -6\ln^3(\frac{\delta}{6Nk^2})) + 1$ .

The above guarantees that if we sample each slot  $K_1$  times the probability that our estimate of the transition probability will be outside our desired bound is less than  $\frac{\delta}{3}$ .

Using the pigeon-hole principle we know that total number of visits to slots marked unknown is  $Nk^2K_1$ . After at most this number of visits all slots will be marked known.

2. The Implicit Exploit or Explore Lemma gives a probability of  $\frac{\alpha}{R_{max}}$  of getting to explore. We now wish to show that after  $K_2$  attempts to explore (i.e. when we do not exploit), we obtain the  $K_1$  required visits. Let  $X_i$  be an indicator random variable which is 1 if we reach to the exploration state ( $G_0$  in Lemma 2) when we do not exploit, and 0 otherwise. Let  $Z_i = X_i - \frac{\alpha}{R_{max}}$ , and let  $Z'_i = \frac{\alpha}{R_{max}} - X_i$ , and apply Chernoff

bound on the sum of  $Z_i$ 's and  $Z'_i$ 's as before. We get that  $\text{Prob}(|\sum_{i=1}^{K_2} X_i - \frac{K_2 \alpha}{R_{max}}| > K_2^{\frac{1}{3}}) < 2e^{-\frac{K_2^{\frac{1}{3}}}{2}}$ . We can now choose  $K_2$  such that  $K_2^{\frac{1}{3}} + K_2 \frac{\alpha}{R_{max}} > k^2 N K_1$  and  $2e^{-\frac{K_2^{\frac{1}{3}}}{2}} < \frac{\delta}{3k^2 N}$  to guarantee that we will have a failure probability of less than  $\frac{\delta}{3}$  due to this reason.

3. When we perform a  $T$ -step iteration without learning our expected return is  $\text{Opt}(\Pi_M(T, \epsilon)) - \epsilon$ . However, the actual return may be lower. This point is handled by the fact that after polynomially many local exploitations are carried out,  $\text{Opt}(\Pi_M(T, \epsilon)) - \frac{1}{2}\epsilon$  can be obtained with a probability of failure of at most  $\frac{\delta}{3}$ . This is obtained by standard Chernoff bounds, and makes use of the fact that the standard deviation of the expected reward in a  $T$ -step policy is bounded because the maximal reward is bounded by  $R_{max}$ . More specifically, consider  $z = MNT$  exploitation stages for some  $M > 0$ . Denote the average return in an exploitation stage by  $\mu$ , and let  $X_i$  denote the return in the  $i$ -th exploitation stage ( $1 \leq i \leq z$ ). Let  $Y_i = \frac{\mu - X_i}{R_{max}}$ . Notice that  $|Y_i| \leq 1$ , and that  $E(Y_i) = 0$ . Chernoff bound implies that:  $\text{Prob}(\sum_{j=1}^z Y_j > z^{\frac{2}{3}}) < e^{-\frac{z^{\frac{1}{3}}}{2}}$ . This implies that the average return along  $z$  iterations is at most  $\frac{R_{max}}{z^{\frac{1}{3}}}$  lower than  $\mu$  with probability of at least  $e^{-\frac{z^{\frac{1}{3}}}{2}}$ . By choosing  $M$  such that  $z > (\frac{2R_{max}}{\epsilon})^3$ , and  $z > 6(\ln(\frac{\delta}{3}))^{-3}$ , we get the desired result: with probability less than  $\frac{\delta}{3}$  the value obtained will not be more than  $\frac{\epsilon}{2}$  lower than the expected value.

By making the failure probability less than  $\frac{\delta}{3}$  for each of the above stages, we are able to obtain a total failure probability of no more than  $\delta$ .

From the proof, we can also observe the bounds on running times required to obtain this result. However, notice that in practice, the only bound that we need to consider when implementing the algorithm is the sample size  $K_1$ .

To remove the assumptions that the  $\epsilon$ -return mixing time is known, we proceed as in (Kearns & Singh, 1998). From the proof of the algorithm we deduced some polynomial  $P$  in the problem parameters such that if  $T$  is the mixing-time, then after  $P(T)$  steps we are guaranteed, with probability  $1 - \delta$ , the desired return. We repeat the execution of the algorithm for all values of  $T = 1, 2, 3, \dots$ , each time performing  $P(T)$  steps. Suppose that  $T_0$  is the mixing time, then after  $\sum_{i=1}^{T_0} P(i) = O(P(T_0)^2)$  steps, we will obtain the desired return. ■

Notice that the R-max algorithm does not have a final halting time and will be applied continuously as long as the agent is functioning in its environment. The only caveat is that at some point our current mixing time candidate  $T$  will be exponential in the actual mixing time  $T_0$ , at which point each step of the algorithm will require an exponential calculation. However, this will occur only after an exponential number of steps. This is true for the  $E^3$  algorithm too.

Another point worth noting is that the agent may never know the values of some of the slots in the game because of the adversary's choices. Consequently, if  $\pi$  is the optimal policy given full information about the game, the agent may actually converge to a policy  $\pi'$  that differs from  $\pi$ , but which yields the best return given the adversary's actual behavior.

This return will be no smaller than the return guaranteed by  $\pi$ . The mixing time of  $\pi'$  will, in general, differ from the mixing time of  $\pi$ . However, we are guaranteed that if  $T_0$  is the  $\epsilon$ -return mixing time of  $\pi$ , and  $v$  is its value, after time polynomial in  $T_0$ , the agent's actual return will be at least  $v$  (subject to the deviations afforded by the theorem).

#### 4.1 Repeated Games

A stochastic game in which the set of stage games contains a single game is called a *repeated game*. This is an important model in game theory and a lot of work has been devoted to the study of learning in repeated games (Fudenberg & Levine, 1993). There is large class of learning problems associated with repeated games, and the problem as a whole is referred to as *repeated games with incomplete information* (Aumann & Maschler, 1995). The particular class of repeated games with incomplete information we are using (i.e., where the agent gets to observe the adversary's actions and the payoffs, and it knows the value of  $R_{max}$ ) is known as an *Adaptive Competitive Decision Process* and has been studied, e.g., by Banos (Banos, 1968) and Megiddo (Megiddo, 1980).

Because a repeated game contains a single stage game, there are no transition probabilities to learn. However, there is still the task of learning to play optimally. In addition, because there is only a single stage game, the mixing time of any policy is 1 – because the agent's expected reward after playing a single stage-game is identical to the policy's expected reward. However, the time required to guarantee this expected reward could be much larger. This stems from the fact that the optimal policy in a game is often mixed. That is, the agent chooses probabilistically, and not deterministically, among different options.

In repeated games, the R-max algorithm is slightly modified, as we do not need to maintain a fictitious state and we need not maintain statistics on the frequency of various transitions. We describe the precise algorithm below:

**Initialization** Initialize the game model with payoffs of  $R_{max}$  for every joint action for the agent and 0 for the adversary. Mark all joint actions as *unknown*.

**Play** Repeat the following process:

**Policy Computation** Compute an optimal policy for the game based on the current model and play it.

**Update** If the joint action played is marked unknown, update the game matrix with its observed payoffs and mark is *known*.

Given  $\epsilon > 0$ , and  $0 < \delta < 1$ , we need to show that after polynomially many iterations  $M$ , where  $M$  is polynomial in the number of entries in the game,  $\frac{1}{\epsilon}$ , and  $\frac{1}{\delta}$ , we obtain a payoff that is at most  $\epsilon$  lower than the expected payoff of the optimal strategy in this game, with probability of at least  $1 - \delta$ .

First notice that the expected payoff at each stage, when we do not expose the value of a new entry, is greater or equal to the expected payoff of the optimal strategy. By choosing  $M = Q_1 + Q_2$ , where  $Q_2 = k^2 R_{max}/M \leq \epsilon/2$  we get that the loss due to learning of new entries is bounded by  $\frac{\epsilon}{2}$ . Now, we need to guarantee that after  $Q_1$  executions (where  $Q_1$  is polynomial in the problem parameters) of a policy with expected payoff greater or equal

to  $r$  (where  $r$  is the expected payoff of the optimal policy in the original game), our actual payoff is at least  $r - \epsilon/2$  with probability of at least  $1 - \delta$ . This follows from the arguments presented in case 3 of the general proof for SGs.

## 5. Conclusion

We described R-max, a simple reinforcement learning algorithm that is guaranteed to lead to polynomial time convergence to near-optimal average reward in zero-sum stochastic games. In fact, R-max guarantees the safety level (probabilistic maximin) value for the agent in general non-cooperative stochastic games.

R-max is an optimistic model-based algorithm that formally justifies the optimism in the face of uncertainty bias. Its analysis is similar, in many respects, to Kearns and Singh’s  $E^3$  algorithm. However, unlike the  $E^3$ , the agent does not need to explicitly contemplate whether to explore or to exploit. In fact, the agent may never learn an optimal policy for the game,<sup>4</sup> or it may play an optimal policy without knowing that it is optimal. The “clever” aspect of the agent’s policy is that it “offers” a catch to the adversary: if the adversary plays well, and leads the agent to low payoffs, then the agent will, with sufficient probability, learn something that will allow it to improve its policy. Eventually, without too many “unpleasant” learning phases, the agent will have obtained enough information to generate an optimal policy.

R-max can be applied to MDPs, repeated games, and SGs. In particular, all single-controller stochastic game instances covered in (Brafman & Tennenholz, 2000) fall into this category, and R-max can be applied to them. However, R-max is much simpler conceptually and easier to implement than the LSG algorithm described there. Moreover, it also attains higher payoff: In LSG the agent must pay an additional multiplicative factor  $\phi$  that does not appear in R-max.

Two other SG learning algorithms appeared in the literature. Littman (Littman, 1994) describes a variant of Q-learning, called minimax Q-learning, designed for 2-person zero-sum stochastic games. That paper presents experimental results, asymptotic convergence results are presented in (Littman & Szepesvri, 1996). Hu and Wellman (Hu & Wellman, 1998) consider a more general framework of multi-agent general-sum games. This framework is more general than the framework treated in this paper which dealt with fixed-sum, two-player games. Hu and Wellman based their algorithm on Q-learning as well. They prove that their algorithm converges to the optimal value (defined, in their case, via the notion of Nash equilibrium). However, convergence is in the limit, i.e., provided that every state and every joint action has been visited infinitely often. Note that an adversary can prevent a learning agent from learning certain aspects of the game indefinitely and that R-max’s polynomial time convergence to optimal payoff is guaranteed even if certain states and joint actions have never been encountered.

The class of repeated games is another sub-class of stochastic games for which R-max is applicable. In repeated games,  $T = 1$ , there are no transition probabilities to learn, and we need not use a fictitious stage-game. Therefore, a much simpler version of R-max can be used. The resulting algorithm is much simpler and much more efficient than previous

---

4. In a game, an agent need not play optimally to obtain an optimal reward because it may obtain this reward because of bad choices by the adversary.

algorithms by Megiddo (Megiddo, 1980) and by Banos (Banos, 1968). Moreover, for these algorithms, only convergence in the limit is proven. A more recent algorithm by Hart and Mas-Colell(Hart & Mas-Colell, 2001) features an algorithm that is much simpler than the algorithms by Banos and Megiddo. Moreover, this algorithm is *Hannan-Consistent* which means that it not only guarantees the agent its safety level, but it also guarantees that the agent will obtain the maximal average reward given the actual strategy used by the adversary. Hence, if the adversary plays sub-optimally, the agent can get an average reward that is higher than its safety-level. However, it is only known that this algorithm converges almost-surely, and its convergence rate is unknown. An interesting open problem is whether a polynomial time hannan-consistent near-optimal algorithm exists for repeated games and for stochastic games.

**Acknowledgments:** We wish to thank Amos Beimel for his help in improving the error bound in Lemma 1 and the anonymous referees for their useful comments. The first author is partially supported by the Paul Ivanier Center for Robotics and Production Management.

## References

- Aumann, R., & Maschler, M. (1995). *Repeated Games with Incomplete Information*. MIT Press.
- Banos, A. (1968). On pseudo games. *The Annals of Mathematical Statistics*, 39, 1932–1945.
- Brafman, R., & Tennenholtz, M. (2000). A near-optimal polynomial time algorithm for learning in certain classes of stochastic games. *Artificial Intelligence*, 121(1–2), 31–47.
- Fudenberg, D., & Levine, D. (1993). Self-confirming equilibrium. *Econometrica*, 61(3), 523–545.
- Hart, S., & Mas-Colell, A. (2001). A general class of adaptive strategies. *Journal of Economic Theory*, forthcoming.
- Hoffman, A., & Karp, R. (1966). On Nonterminating Stochastic Games. *Management Science*, 12(5), 359–370.
- Hu, J., & Wellman, M. (1998). Multi-agent reinforcement learning: Theoretical framework and an algorithms. In *Proc. 15th ICML*.
- Kaelbling, L. P. (1993). *Learning in Embedded Systems*. The MIT Press.
- Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement learning: A survey. *Journal of AI Research*, 4, 237–285.
- Kearns, M., & Koller, D. (1999). Efficient reinforcement learning in factored mdps. In *Proc. 16th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 740–747.
- Kearns, M., & Singh, S. (1998). Near-optimal reinforcement learning in polynomial time. In *Int. Conf. on Machine Learning*.

- Littman, M. L. (1994). Markov games as a framework for multi-agent reinforcement learning. In *Proc. 11th ICML*, pp. 157–163.
- Littman, M. L., & Szepesvri, C. (1996). A generalized reinforcement-learning model: Convergence and applications. In *Proc. 13th Intl. Conf. on Machine Learning*, pp. 310–318.
- Megiddo, N. (1980). On repeated games with incomplete information played by non-bayesian players. *Int. J. of Game Theory*, 9, 157–167.
- Monderer, D., & Tennenholtz, M. (1997). Dynamic Non-Bayesian Decision-Making. *J. of AI Research*, 7, 231–248.
- Moore, A. W., & Atkenson, C. G. (1993). Prioritized sweeping: Reinforcement learning with less data and less real time. *Machine Learning*, 13.
- Schmidhuber, J. H. (1991). Curious model-building control systems. In *Proc. Intl. Joint Conf. on Neural Networks*, pp. 1458–1463.
- Sutton, R. S. (1990). Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *Proc. of the 7th Intl. Conf. on Machine Learning*. Morgan Kaufmann.
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. MIT Press.
- Tadepalli, P., & Ok, D. (1998). Model-based average reward reinforcement learning. *Artificial Intelligence*, 100, 177–224.

# Bidding Clubs: Institutionalized Collusion in Auctions

Kevin Leyton-Brown  
 Dept. of Computer Science  
 Stanford University  
 Stanford, CA 94305  
 kevinlb@stanford.edu

Yoav Shoham  
 Dept. of Computer Science  
 Stanford University  
 Stanford, CA 94305  
 shoham@cs.stanford.edu

Moshe Tennenholtz  
 Dept. of Computer Science  
 Stanford University  
 Stanford, CA 94305  
 moshe@robotics.stanford.edu

## ABSTRACT

We introduce a class of mechanisms, called *bidding clubs*, for agents to coordinate their bidding in auctions. In a bidding club agents first conduct a “pre-auction” within the club; depending on the outcome of the pre-auction some subset of the members of the club bid in the primary auction in a prescribed way; and, in some cases, certain monetary transfers take place after the auction. Bidding clubs have self-enforcing collusion properties in the context of second-price auctions. We show that this is still true when multiple auctions take place for substitutable goods, as well as for complementary goods. We also present a bidding club protocol for first-price auctions. Finally, we show cases where bidding clubs have self-enforcing cooperation protocols in arbitrary mechanisms.<sup>1</sup>

## 1. INTRODUCTION

With the exploding popularity of auctions on the Internet and elsewhere has come increased interest in systems to assist (software or human) agents bidding in such auctions. Most of these systems have to date done little more than aggregate information from multiple auctions and present it to the user in a convenient fashion (e.g., [www.auctionwatch.com](http://www.auctionwatch.com)). There is now beginning to emerge a second generation of systems which actually provide bidding advice and automation services to bidders, going beyond the familiar proxy-bidding feature prevalent in online auctions to the realm of bona-fide decision support.

This paper looks even beyond such systems, which are geared towards assisting a single bidder, and presents a class of systems to assist a collection of bidders, “bidding clubs”. The idea is similar to the idea behind “buyer clubs” on the Internet (e.g., [www.merkata.com](http://www.merkata.com) and [www.mobshop.com](http://www.mobshop.com)), namely to aggregate the market power of individual bidders. The new twist is that whereas in a buyer club there is a perfect alignment of the various buyers’ interests (since there

<sup>1</sup>This work was partly supported by DARPA grant number F30602-98-C-0214-P00005.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC’00, October 17-20, 2000, Minneapolis, Minnesota.  
 Copyright 2000 ACM 1-58113-272-7/00/0010 ..\$5.00

the more buyers join in a purchase the lower the price for everyone), in a bidding club there is a more complex strategic relationship among them, and the bidding club rules must be designed accordingly.

Here’s a simple example. Consider an auction with a single seller, and six potential buyers. Assume that three of the potential buyers – *A*, *B* and *C*, with corresponding (secret) valuations  $v_1 > v_2 > v_3$  – attempt to coordinate their bidding. Assume the auction is a first-price auction. Under well known assumptions from the auction literature, it would be the interest of each bidder to bid exactly  $5/6$  of his true value in the auction. Thus *A* would end up with a surplus of  $v_1/6$  (if he wins the auction) or 0 (if he doesn’t), and *B* and *C* with a surplus of 0. Is there some pre-agreement *A*, *B* and *C* can make that will cause all of them to come out of the auction at least as well off, and some of them strictly better off? One could naïvely say that they would each reveal their valuations to one another agreeing that only the highest would go on to the auction; *A* would therefore be the one going on, and when he bids in the auction he would bid lower than  $5v_1/6$  (a bid of  $3v_1/4$  will work, given the above-mentioned assumptions), and thus increase his expected surplus. The obvious flaw in this mechanism is that *A*, *B* and *C* will have incentive to lie in this initial phase; this could still be true if *A* were obliged to pay *B* and *C* a certain amount if they sat it out and he won the auction.

The above protocol is a simple instance of the class *bidding clubs*. In general, given some primary mechanism (typically, an auction), a bidding club protocol is as follows:

1. Some set of bidders are invited to join the bidding club, and informed of its rules. The other bidders are not made aware of the existence of the bidding club; we assume here that they are not even aware of the possibility of its existence.
2. The bidders have the freedom to join the club or not. If they do it is assumed that they are guaranteed to follow its rules.<sup>2</sup>
3. The bidding-club coordinator (or simply ‘coordinator’) asks the members for certain private information, such as their valuations for the good that is being sold. Notice that in general bidders may cheat about their valuations.

<sup>2</sup>In practice, we will design bidding clubs in such a way that any agent who would want to participate in the main auction will want to join the bidding club.

4. The coordinator determines, according to pre-specified rules, how the members should behave in the primary mechanism based on the information they all supply.
5. The coordinator may also determine (and enforce) additional monetary transfers of the club members, based on the results of the main mechanism.
6. The coordinator acts only as a representative of bidders.

It may seem natural to ask why a coordinator should be willing and/or able to function as a trusted third party, without attention having been paid to its own incentives. We believe that it is best not to see the coordinator as a party (with interests of its own) at all; rather, we conceive of a coordinator as a software agent which is able to act only according to its (commonly-known) programming. It is therefore possible for the coordinator to act reliably—and for agents to be confident that the coordinator will act reliably—even in cases where the coordinator stands to gain nothing through its efforts. We do assume that coordinators should not cost money to operate—all of our coordinators are budget-balanced except for one that (unavoidably!) makes money. Finally, we have often been asked about the legal issues surrounding the use of bidding clubs. While this is an interesting and pertinent question, it exceeds both our expertise and the scope of this paper.

It turns out that, while the simple mechanism outlined earlier fails, a more sophisticated one will ensure that B and C do not participate in the primary auction, and that A is therefore assured higher expected payoff in the auction. More generally, the contributions of this paper are as follows:

1. We present a protocol for self-enforcing cooperation in second-price auctions for substitute goods.
2. We present a protocol for self-enforcing cooperation in second-price auctions for complementary goods.
3. We present a protocol for self-enforcing collusion in first-price (as well as Dutch) auctions, in which only some of the agents coordinate their activities, and which does not make any use of monetary transfers.
4. We present a protocol for self-enforcing cooperation in general auctions and economic mechanisms, when the agents' types (e.g. valuations for goods) are taken from a finite set.

## 2. TECHNICAL BACKGROUND

The strategic interaction among self-interested agents is a primary topic of study in microeconomics [4] and game theory [1]. In particular, the design of protocols for strategic interactions is the subject of the field termed *mechanism design* [1]. The role of a mechanism (in particular, auction) designer is to define a game whose equilibrium strategies are desirable in some respect or another. Thus, the design of a bidding club consists of taking a given mechanism – the primary auction – and turning it into a more elaborate one, namely one with an added first stage in which a subset of the players play in some newly-designed game (as well as some additional rules regarding behavior in the primary auction and possible side payments after the auction).

Research on strategic aspects of multi-agent activity in Artificial Intelligence has grown rapidly in the recent years. This work has concentrated on the design of protocols for agents' interaction [7, 3, 9], and shares much in common with work on mechanism design in economics. Many principles and ideas grew up from the mechanism design literature, and have been adapted to the AI context.

Although the study of deals among agents has received much attention in the AI literature (see e.g. [7]), and although the study and design of contracts is central to information economics [4] (and received much attention in the recent AI literature [8]), the literature on cooperation under incomplete information in auctions and trades is quite limited. In particular, the literature on collusion in auctions is somewhat spotty. It is still too broad to give a complete overview of it, and the bulk of it is informal. In the formal literature on the topic, the results are quite specific, and certainly do not apply in settings of parallel auctions (with either substitutability or complementarity among goods), first-price auctions without side-payments, and general mechanisms, which are the focus of our technical results. The closest result from the literature of which we are aware is by Graham and Marshall [2], who present a protocol for self-enforcing collusion by a subset of the participants of a (single-good) second-price auction. We discuss this result below. Additional related study of collusion in auctions can be found in [5].

## 3. AUCTION PRELIMINARIES

We now present some preliminaries of auction theory, as well as a description of the classical auction model discussed in the paper and our parallel auction model.

### 3.1 Single auctions

An auction procedure for selling a single good to one of  $n$  potential participants,  $N = \{1, 2, \dots, n\}$  is characterized by 4 parameters,  $M, g, c, d$ :  $M$  is the set of possible messages a participant may submit;  $g = (g_1, g_2, \dots, g_n)$ ,  $g_i : M^n \rightarrow [0, 1]$ , is an allocation function, where  $g_i$  determines the probability the winner of the auction will be agent  $i$ ;  $c : M^n \rightarrow R$  determines the payment by the winner of the auction;  $d$  is a participation fee. It is assumed that agents may decide not to participate in an auction.

In order to analyze auctions we have to discuss the information available to the participants. We assume the independent private values model, with no externalities. Each agent  $i$  is assumed to have a valuation  $v_i$  selected from the interval of real numbers  $[0, 1]$  or from a finite domain, which captures its maximal willingness to pay for the good. We further assume that this valuation is selected from the uniform distribution on the interval  $[0, 1]$  or on a finite domain. For ease of presentation we will assume the continuous case, excluding the section on general mechanisms, where the assumption that the set of possible valuations is finite is required for our result. If agent  $i$  obtains the good and is asked to pay  $p$ , as well as a participation fee  $d$ , then its utility,  $u_i$ , is given by  $v_i - p - d$ ; otherwise, if it is not assigned any good then its utility is  $-d$ ; if the agent does not participate in the auction then its utility is 0.

The above defines a Bayesian game, where a strategy for an agent is a decision about the message to be sent given its valuation, and the payoffs are determined as above. The solution of this game is given by computing a (Bayesian Nash)

equilibrium of it: a joint strategy of the agents such that it is irrational for each agent to deviate from its strategy, given that all of the other agents stick to their strategy. Given an equilibrium strategy  $b = (b_1, b_2, \dots, b_n)$ , one can compute  $L_i(b)$ , the expected utility of agent  $i$  in equilibrium of the corresponding game. In a case where there is more than one equilibrium  $L_i(b)$  is taken as the lowest expected utility over all the equilibria. Further discussion of equilibrium uniqueness is omitted from this paper.

One of the best-known auction mechanisms is the second-price auction. In such an auction, each participant submits a bid in a sealed envelope. The agent with the highest bid wins the good and pays the amount of the second-highest bid, and all other participants pay nothing. In a case of a tie, the winner of the auction is selected randomly, with uniform probability. If there is no participation fee then participation in second-price auctions is always rational. Truth revealing, i.e.  $b_i(v_i) = v_i$ , is an equilibrium of the second-price auction (in fact, it is an equilibrium in dominant strategies). Another popular auction is the first-price auction. These auctions are conducted similarly to second-price auctions, except that the winner pays the amount of his own bid. The equilibrium analysis of first-price auctions is quite standard. For example, if valuations are selected according to the uniform distribution on  $[0, 1]$  and there is no participation fee, then the strategy of agent  $i$  in equilibrium is  $b_i(v_i) = \frac{n-1}{n}v_i$ .

### 3.2 Parallel auctions

More generally, several auctions may be conducted in parallel. We first consider the case of two parallel auctions of similar goods. A parallel auction is given in this case by a pair  $A = (A_1, A_2)$ , where  $A_i = \langle N, g, c, d \rangle$ , ( $i = 1, 2$ ) as before.

One such problem is a *parallel auction for substitute goods*, in which the set of possible buyers  $N$  is shared among  $A_1$  and  $A_2$ , and each agent's valuation for the pair of goods  $\{g_1, g_2\}$  equals its valuation for  $g_1$  which equals its valuation for  $g_2$ . Agent  $i$ 's strategy consists of two parts:

1. It selects at most one of the auctions, in which it will participate.
2. It submits a bid in the selected auction.

Parallel auctions for substitute goods define a Bayesian game in a natural way. For example, if the auctions are second-price auctions, then an appropriate equilibrium of the corresponding parallel auction is as follows: each agent randomly selects one of the auctions, and sends his actual valuation as his bid there.

Another type of parallel auction is the *parallel auction for complementary goods*. Here we have two similar auctions, e.g. second-price auctions, for two different goods  $g_1$  and  $g_2$ . The set of agents  $N = N_1 \cup N_2 \cup N_p$  consists of three parts:

- $N_1$  are agents that are interested only in  $g_1$
- $N_2$  are agents that are interested only in  $g_2$
- $N_p$  are agents that have valuation 0 for  $g_1$  and for  $g_2$ , but their valuation for the pair  $\{g_1, g_2\}$  is uniformly distributed on the interval  $[0, 2]$ .

For ease of exposition we will assume that we can distinguish whether an agent is from group  $N_1, N_2$ , or  $N_p$ , and that the agents in  $N_p$  have extremely high negative utility for losses. This second assumption means that an agent will never submit bids in both auctions; notice that we assumed that an agent who is interested in obtaining a pair of goods has a valuation of 0 for getting only one of them, and therefore by bidding in two auctions the agent may end up getting and paying for only one good. Hence, we will assume that the strategies available to the agents are as in the case of substitute goods.

We will rely on the notion of surplus in our evaluation of coordinators for parallel auctions. The surplus of an allocation is defined as the sum of agents' valuations for that allocation. For example, in a parallel auction for substitute goods the surplus of an allocation that assigns good  $g_1$  in auction 1 to agent  $i$ , and assigns good  $g_2$  in auction 2 to agent  $j$ , is  $v_1(g_1) + v_2(g_2)$  (i.e., the sum of these agents' valuations for the goods they are assigned).

## 4. COORDINATORS AND BIDDING CLUBS

Let  $G \subset N$ , where  $1 < |G| < n$ . W.l.o.g let the elements of  $G$  be  $\{1, 2, \dots, |G|\}$ . Given an auction  $A$ , denote by  $\Phi_i(A)$  ( $1 \leq i \leq n$ ) the set of strategies available to agent  $i \in N$ .

Given a set of *coordinator messages*,  $M_c$ , which we take w.l.o.g to be  $R_+$ , a (*bidding club*) *coordinator* is a pair of functions  $C(A, G) = (T_1(A, G), T_2(A, G))$ , where  $T_1(A, G) : M_c^{|G|} \rightarrow \Phi_i(A)^{|G|}$  and  $T_2(A, G) = (t_1^c, t_2^c, \dots, t_{|G|}^c)$ ,  $t_i^c : M_c^{|G|} \times M^n \rightarrow R$ . Namely, a coordinator is a mechanism that asks the agents in  $G$  for some information and decides on the way they will behave in  $A$ ; this is determined by the function  $T_1(A, G)$ . In addition, following the decision made by  $T_1(A, G)$ , and given the messages sent in the main auction  $A$  by members of  $N \setminus G$ , an additional payment  $t_i^c$  may be imposed on agent  $i$ . The payment can be negative, positive, or zero.  $M_c$  contains the null message  $e$  that tells the coordinator that the corresponding agent is not willing to participate in the coordination activity. This agent will be free to participate in the auction by itself, and will not be asked to make any payments to the coordinator. A key assumption is that participants in  $N \setminus G$  are unaware of even the possibility of the existence of a coordinator, and that they act according to an equilibrium of  $A$ . We denote the game obtained by concatenating  $C(A, G)$  and  $A$ , by  $\bar{C}(A, G)$ . For every agent  $i$ , let  $L_i(A)$  be the agent's expected utility in an equilibrium of  $A$ , and let  $L_i(\bar{C}(A, G))$  be the agent's expected utility in an equilibrium of  $\bar{C}(A, G)$ .

**DEFINITION 1.** *Given an auction  $A$ , and a  $G \subset N$  as before, we will say that a participation-preserving coordinator for  $G$  in  $A$  exists, if there exists  $C(A, G)$ , such that every agent  $i \in G$  that would have had participated in  $A$  will also participate in  $C(A, G)$  (in equilibrium of  $\bar{C}(A, G)$ ).*

**DEFINITION 2.** *We say that a utility-improving coordinator exists if there exists a participation-preserving coordinator, and  $L_i(\bar{C}(A, G)) > L_i(A)$  (i.e. participation in the bidding club is beneficial).*

The existence of a utility-improving coordinator for an auction setup implies a self-enforcing cooperative strategy for a group of agents.

**DEFINITION 3.** We say that a surplus-improving coordinator for  $G$  in  $A$  exists if there exists a  $C(A, G)$  that is participation-preserving, and the expected surplus of the members of  $G$  in  $\bar{C}(A, G)$  is greater than their expected surplus in  $A$ .

When dealing with parallel auctions in sections 5.2 and 5.3, we will be interested in surplus-improving coordinators. Besides the observation that neither concept implies the other, the discussion of the connection between utility-improving and surplus-improving coordinators is left to the full paper.

## 5. COORDINATION IN SECOND-PRICE AUCTIONS

### 5.1 Second-price auctions for a single good

The case of collusion in second-price auctions is discussed in [2]. The following theorem may be deduced from this work; we present the result here for the sake of completeness. Consider a second-price auction. In the case of a second-price auction a group of buyers may wish to avoid paying a participation fee, or alternatively bidders who will certainly lose may want to receive advance notice. As it turns out, such behavior can be obtained:

**THEOREM 1.** There exists a utility-improving coordinator for second-price auctions.

#### Sketch of proof:

In the case of a second-price auction, no assumptions on the distribution of the agents' valuations need to be made. We will assume that there is a participation fee  $d > 0$ , and show a coordination protocol that enables the members of the group  $G$  who do not have the highest valuation to avoid paying  $d$ . We use the following protocol:

1. The agents in  $G$  are asked to submit their valuations to the coordinator.
2. Let  $v_1$  and  $v_2$  denote the highest and second highest valuations, announced by agents 1 and 2, respectively.<sup>3</sup>
3. Only agent 1 is represented in the main auction, and his bid there will be  $v_1$ .
4. If agent 1 wins the main auction, and is asked to pay  $z$ , and  $z < v_2$ , then agent 1 will pay  $v_2 - z$  to the coordinator.

We show that if the agents participate in the pre-auction and reveal their true valuations there, then this cooperation will be beneficial to them. The agent with the highest valuation cannot lose, because his behavior and expected gain will be as in the case where there was no coordinator. The other agents will gain due to the fact they won't need to pay the participation fee.

Consider now the agent  $i \in G$  with the highest valuation, and assume that the other agents in  $G$  are truth-revealing agents. Given that truth-revealing is an equilibrium of second-price auctions, agents in  $N \setminus G$  are taken to

---

<sup>3</sup>Note that, unlike in some of the coordination protocols that follow, the coordinator behaves the same regardless of whether some bidders decline to participate in the coordination.

be truth-revealing as well. Given that if the agent  $i$  wins the main auction, then he pays exactly the highest valuation in  $N - \{i\}$  (because he will pay the maximum of the auction's second-highest bid and  $v_2$ ). Standard second-price auction analysis yields that it is irrational for  $i$  to deviate from truth-revealing to the announcement of a higher valuation. If agent  $i$  was willing to participate in the main auction then clearly he does not wish to lose the pre-auction and therefore announcing a lower valuation than his actual one is irrational too. Clearly, every agent  $j \neq i, j \in G$  does not have any incentive to cheat if the others are truth-revealing. He can only lose if by cheating he will be chosen to participate in the main auction. ■

It is easy to see that our result holds for Japanese auctions as well. In a Japanese auction an auctioneer starts with a low asking price, and continuously increments this price as long as are still multiple agents willing to pay the current price. Once only a single agent remains, he will get the good for the current asking price. The fact our result holds also for Japanese auctions is immediately implied by the fact that in both Japanese auctions and second-price auctions the good is sold to the agent with the highest valuation, at a price that equals the second-highest valuation.

### 5.2 Parallel auctions with substitute goods

In this section we deal with parallel auctions of substitute goods. Here the idea of the coordinator is to ensure that the two agents with the highest valuations in the group  $G$  will compete for different goods rather than among themselves. This will enable to improve upon the surplus of the members of  $G$ . We can show:

**THEOREM 2.** There exists a surplus-improving coordinator for parallel second-price auctions of substitute goods.

#### Sketch of proof:

1. The agents in  $G$  are asked to submit their valuations to the coordinator.
2. Let  $v_1, v_2$ , and  $v_3$  denote the highest, the second highest, and the third highest valuations which have been announced, respectively.<sup>4</sup>
3. Only the agents with the highest and second highest valuations will participate in the main auction. The agents will be randomly assigned to different auctions.
4. If an agent gets the object in auction  $A_i$  for the price  $y < v_3$ , then he will pay  $v_3 - y$  to the coordinator.

It is clear that if all agents obey the coordinator's protocol, and send their actual valuations to the coordinator, then the agents will improve upon their surplus. In equilibrium agents will want to participate; for example, consider agents 1 and 2, having the two highest bids submitted to the coordinator. As a result of the coordination the first agent will have a lower expected payment, since he will always pay some amount less than  $v_2$ , while the second agent will have a greater chance of winning, since he will never be outbid by agent 1.

---

<sup>4</sup>Once again, note that the coordinator behaves the same regardless of whether some bidders decline to participate.

We now show that truth-revealing is an equilibrium. Consider an agent  $i_1$ , with the highest valuation in  $G$ ,  $v_1$ , and assume that the rest of the agents are truth-revealing. If agent 1 reports a valuation higher than  $v_1$ , and obtains as a result of this a good he could not obtain otherwise, then it must be the case that his payment is higher than his valuation, which makes that deviation irrational. It is clear that reporting on a valuation lower than  $v_1$  does not help agent 1.

Consider an agent  $i_2$ , with the second-highest valuation in  $G$ ,  $v_2$ , and assume the other agents are truth-revealing. If the agent reports a higher valuation than  $v_1$  then he will be the highest-ranking bidder in the pre-auction rather than the second highest-ranking, but this will not benefit him as the top two bidders are assigned to auctions randomly. The rest of the analysis is the same as for  $i_1$ .

Consider an agent  $i_3$ , with the third-highest valuation in  $G$ ,  $v_3$ , and assume the other agents are truth-revealing. If the agent reports a valuation that causes it to gain the pre-auction, then its payment will be at least  $v_2 > v_3$ , which makes such deviation irrational. Similar analysis will hold for agents with lower valuations. ■

### 5.3 Parallel auctions with complementary goods

In this section we deal with parallel auctions for complementary goods. Our aim is to allow the participants in  $G$  to obtain a higher surplus than what they could obtain without the coordinator. We assume that in  $G$  we have at least two representatives of  $N_1, N_2$  and  $N_p$ . We can show:

**THEOREM 3.** *There exists a surplus-improving coordinator for parallel second-price auctions of complementary goods.*

#### Sketch of proof:

Let  $0 < k \ll 1$  be a commonly-known constant. We will use the following coordinator<sup>5</sup>:

1. The coordinator asks the agents that are interested in the single goods for their valuations
2. The coordinator selects two agents,  $s_1$  and  $s_2$ , who reported the highest valuations for goods  $g_1$  and  $g_2$ ,  $v_1$  and  $v_2$  respectively.
3. If any agent from  $N_1 \cup N_2$  declined to participate, the coordinator submits bids in the appropriate auctions for all agents in  $N_1 \cup N_2$  who did elect to participate, with a price offer equal to the agents' stated valuations, and the protocol is complete. Otherwise, if all agents elected to participate, we proceed to step 4.
4. The coordinator announces  $v_1$  and  $v_2$  to all of the participants in  $G$ .
5. The coordinator asks the agents that are interested in the pair of goods for their valuations.
6. The coordinator randomly selects an agent,  $s_p$ , who reported a valuation  $v_p$  for the pair of goods, such that  $v_1 + v_2 + 2k < v_p$  (if such an agent exists).

<sup>5</sup>This requires a quite straightforward modification to the definition of coordinators, which we skip. Namely, a coordinator can run a multi-stage game instead of the function  $T_1(A, G)$ .

7. The coordinator bids  $v_1$  in  $A_1$ , and  $v_2$  in  $A_2$ .
8. If the coordinator wins both auctions, and an agent  $s_p$  exists, then  $s_p$  will get the pair of goods and pay  $v_{sec_1} + v_{sec_2}$  to the coordinator, where  $v_{sec_i}$  is the second-highest bid in  $A_i$ . Agent  $s_p$  will also pay agent  $i$  ( $i = 1, 2$ )  $k + \max(0, v_i - v_{sec_i})$ .
9. If the coordinator only wins auction  $i$ , or if the coordinator wins both auctions but there does not exist an agent  $s_p$ , then agent  $s_i$  gets the good and pays  $v_{sec_i}$  to the coordinator.

Consider an equilibrium of the corresponding  $\bar{C}(A, G)$ , and an agent  $s'_i \in N_i \cap G$  ( $i = 1, 2$ ). It is clear that in equilibrium  $s'_i$  will participate in  $C(A, G)$  and that the submission of a valuation which is at least as high as  $s'_i$ 's valuation by  $s'_i$  dominates the submission of a lower valuation. This is due to the fact that by submitting a valuation that is lower than his actual valuation an agent can only lose, given that this is a second-price auction. The agent cannot lose by participating in the pre-auction, since it is guaranteed to get at least the difference between its stated valuation and the second-highest bid, if its stated valuation is the highest. Moreover, if agent  $s_p$  wins the good then  $s'_i$  may also get a payment of  $k > 0$ . For this reason, and also because  $v_{sec_i}$  may be less than the highest rejected bid from  $N_i \cap G$ , truth revelation will not be in the best interest of agent  $s'_i$ . Instead, he will submit a bid that exceeds his true valuation.

Given the above, an agent  $s_p$ , who has interest in the pair of goods will be willing to participate in the coordinator's protocol if  $v_1 + v_2 + 2k < v_p$ . Note that all agents are aware of  $k$  before placing their bids. It is easy to check that it is irrational for  $s_p$  to send a message that could win the pre-auction if its valuation is smaller than  $v_1 + v_2 + 2k$ , and likewise it is irrational for  $s_p$  to falsely submit a valuation smaller than  $v_1 + v_2 + 2k$ . Otherwise the amount submitted by  $s_p$  is irrelevant, as the coordinator chooses randomly between eligible agents in  $N_p$ . Thus, expected surplus is increased by this protocol. ■

## 6. COORDINATION IN FIRST-PRICE AUCTIONS

**THEOREM 4.** *There exists a utility-improving coordinator for first-price auctions.*

#### Sketch of proof:

Recall that we assume that the agents' valuations are drawn uniformly from the interval  $[0, 1]$ . Our protocol can be easily modified to deal with other distributions on the agents' types. Let  $m$  be the number of agents who will participate in the main auction, who are not members of the bidding club (and who are thus assumed not to be aware even of the possibility of its existence). We use the following protocol:

1. Invite the agents in  $G$  to submit their valuations to the coordinator.
2. If any agent declines to participate, submit bids for all agents that did elect to participate, with a price offer of  $\frac{n-1}{n}v_i$ , and the protocol is complete. Otherwise, if all agents elected to participate, we proceed to step 3.

3. Let the two agents with the highest reported valuations be agents 1 and 2, with reported valuations  $v_1$  and  $v_2$  respectively.
4. If  $\frac{v_1^n}{n} < v_2^m \cdot (v_1 - v_2)$ , submit a bid only for agent 1, with a price offer of  $v_2$ .
5. Otherwise, submit bids for all agents  $i \in G$ , with price offer  $\frac{n-1}{n} v_i$ .

First, we show that if the agents reveal their true valuations then beneficial cooperation ensues. It is clear that the only agent who can gain is the agent with the highest valuation,  $v_1$ , while the other agents do not lose. Note that  $\frac{v_1^n}{n}$  is the expected utility of agent 1 at the equilibrium in the original mechanism, while  $v_2^m \cdot (v_1 - v_2)$  is his expected utility if he submits a bid of  $v_2$  in a modified mechanism with  $m+1$  participants.  $v_1$  benefits because the protocol is tailored specifically to him: the coordinator offers agent 1 the choice of participating in the original mechanism at its equilibrium, or of eliminating some bidders from the auction and bidding  $v_2$ . In every situation, the coordinator selects the alternative that agent 1 would prefer, given his stated valuation. (Note that there exists a set with non-zero measure of values of  $v_1$  and  $v_2$  satisfying the condition in step 3 of the protocol; the demonstration of this fact is left to the full version of the paper.) At the same time, no bidder suffers from being eliminated: each eliminated bidder is assured that a bid will be placed in the main auction exceeding his valuation.

Now we show that the protocol leads the agents to reveal their true valuations. As a result, participation will be rational for all agents. To show that truth-revelation is an equilibrium, assume that all but one of the agents submit their true valuations. Notice that since only agent 1 can profit from the bidding club, the only reason that any agent other than agent 1 would lie is to become the agent with the highest valuation. However, this agent would then either be represented in the original mechanism above the equilibrium, or be made to bid  $v_1$ , more than his valuation. Agent 1 has no reason to lie because the mechanism is tailored exactly to him, as described above. ■

Note that, paradoxically, the bidding club can also benefit bidders who don't even know of its existence! This is due to the fact that in equilibrium of first-price auctions, bids are decreasing as a function of the number of participants, and we assume that all agents are made aware of the number of bidders participating in the main auction.<sup>6</sup> Bidders who are unaware of the bidding club will thus submit lower bids if the bidding club eliminates bidders than if it does not. We do not analyze the case where bidders who are unaware of the bidding club are aware of the total number of bidders including those eliminated by the coordinator, since this knowledge would lead them to knowledge of the bidding club's existence (when they observed that a smaller number of bids were actually entered in the auction), violating a key assumption of our model.

---

<sup>6</sup>We assume that the number of bidders participating in the auction is determined according to the number of distinct bidders wanting to submit bids. Thus if the coordinator places only one bid in the main auction then bidders who are unaware of the bidding club will also be unaware of bidders who were eliminated in the bidding club's pre-auction.

It is easy to see that our result holds for Dutch auctions as well. In a Dutch auction the auctioneer starts with a high asking price, and then continuously decrements this price until an agent claims the good for the current price. The fact our result holds also for Dutch auctions is immediately implied by the strategic equivalence between first-price auctions and Dutch auctions.

## 7. BIDDING CLUBS FOR GENERAL MECHANISMS

The first-price and the second price auctions are two representative auctions, but many other auctions, as well as other economic mechanisms (various types of trades, negotiations, etc.), are also discussed in the literature. In this section we show that utility-improving coordinators exist for many other related contexts as well.

General mechanisms are usually analyzed using Bayesian games. In a Bayesian game each agent has a set of possible types, and an agent's strategy is a decision of his action as a function of his type. The actual type of the agent is known to him, and is selected from a commonly known distribution function. The payoff of each agent is a function of both the joint strategy of the agents and the particular type of the agent. In the context of auctions, the types of the agents refer to their valuations. The definition and analysis of equilibrium strategies for general mechanisms will therefore be similar to what we described in Section 3 for the case of auctions.

In order to prove results that are general and hold for any mechanism, researchers have used the following observation, which is a direct implication of the definition of an equilibrium of a Bayesian game. It turns out that it is enough to consider only mechanisms such that in the equilibrium of the corresponding Bayesian game the agents will reveal their true types. According to this observation, termed the *revelation principle*, it is natural to restrict our attention to (main) mechanisms which make a decision based on true information supplied by the agents.

This brings us to the following general problem. Assume that the agents' types are selected from a finite set, and that the agents are about to participate in a given truth revealing mechanism  $M$ . Assume that the equilibrium of the game associated with that mechanism leads to a non Pareto-optimal outcome for at least one tuple of agent types (i.e. for this tuple of types the agents would better perform a joint strategy that is different from the equilibrium strategy). Can a coordinator be used in order to make a cooperative (beneficial and incentive compatible) deal among the agents? In the sequel, we assume that the valuations of the agents are taken from  $V = \{v_1, \dots, v_m\}$  where  $v_i < v_{i+1}$  for every  $i$ . We can show:

**THEOREM 5.** *Consider a truth revealing mechanism with unique strict Bayesian equilibrium, that leads to a non Pareto-optimal outcome for at least one tuple of agent types. Then, a utility-improving coordinator exists.*

**Basic idea behind proof:** Each agent will be invited to send his valuation to the coordinator. The coordinator will calculate a tuple of other valuations that would benefit the agents (assuming they reported their actual valuations), if submitted to the main mechanism. Notice that while an

agent would lose in equilibrium by deviating from truth-revelation in the original mechanism, sending true valuations is not necessarily an equilibrium if the coordinator submits the new tuple. However, we can show that there exists a useful coordinator which also maintains incentive compatibility.

1. Invite the agents to submit their valuations to the coordinator.
2. If any agent declines to participate, submit the declared valuations of all participating agents to the main mechanism.
3. Otherwise, submit the new tuple of valuations to the main mechanism on behalf of all agents with probability  $p$ ; with probability  $1 - p$  submit the valuations reported by the agents.

The probability  $p$  is determined as follows. Consider an agent  $i$ , who made the announcement  $v_i$ . First, we can compute the maximum expected gain,  $g_i$ , that  $i$  could achieve by submitting a valuation  $v'_i \neq v_i$ . Second, we can compute  $i$ 's smallest expected loss in the original mechanism,  $l_i$ , if  $v_i$  is a false valuation. Notice that  $l_i$  is positive, given the assumption that truth-revelation is a strict Nash equilibrium. Let  $g = \max_i(g_i)$  and  $l = \min_i(l_i)$ . Then we can take  $p = \frac{l}{g+l}$ .

The analysis of this protocol is straightforward. Agents should want to participate, as their expected utility is increased. Incentive compatibility is ensured because the most an agent can gain by lying is  $p \cdot g - (1-p) \cdot l = \frac{l}{g+l} \cdot g - \frac{g}{g+l} \cdot l = 0$ . On expectation agents will lose by lying, since  $g$  and  $l$  are calculated globally, not individually for each agent. ■

## 8. CONCLUSION

In this paper we have presented the notion of bidding clubs and its use in obtaining self-enforcing cooperation in classical auction setups. We have presented protocols for parallel second-price auctions for substitutable and complimentary goods, for first-price auctions for single goods, and for general mechanisms under various assumptions. Our work can be considered as a first attempt to formalize “strategic buyers’ clubs”, where participants may cheat about their valuations and so the club’s protocol must be designed carefully enough to account for this possibility. The study of bidding clubs is complementary to the rich work on efficient market design [4, 1, 6]. Bidding clubs take the agents’ perspective in improving their situation in existing markets, rather than taking a center’s perspective on optimal, revenue maximizing market design.

## 9. REFERENCES

- [1] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.
- [2] D.A. Graham and R.C. Marshall. Collusive bidder behavior at single-object second-price and english auctions. *Journal of Political Economy*, 95:579–599, 1987.
- [3] S. Kraus. Negotiation and cooperation in multi-agent environments. *Artificial Intelligence*, 94, 1997.
- [4] D. Kreps. *A Course in Microeconomic Theory*. Princeton University Press, 1990.
- [5] R.P. McAfee and J. McMillan. Bidding rings. *The American Economic Theory*, 82:579–599, 1992.
- [6] D. Monderer and M. Tennenholtz. Optimal Auctions Revisited. *Artificial Intelligence*, forthcoming, 2000.
- [7] Jeffrey S. Rosenschein and Gilad Zlotkin. *Rules of Encounter*. MIT Press, 1994.
- [8] T. Sandholm and V. Lesser. Leveled commitment contracts. Working paper (to appear in Games and Economic Behavior), 2000.
- [9] M. Tennenholtz. Electronic commerce: From game-theoretic and economic models to working protocols. In *IJCAI-99*, 1999.

## APPENDIX F

# Bidding Clubs in First-Price Auctions

Kevin Leyton-Brown, Yoav Shoham and Moshe Tennenholtz  
Department of Computer Science  
Stanford University  
Stanford, CA 94305  
Email: {kevinlb;shoham;moshe}@cs.stanford.edu

Version: March 22, 2002

We introduce a class of mechanisms, called *bidding clubs*, that allow agents to coordinate their bidding in auctions. Bidding clubs invite a set of agents to join, and each invited agent freely chooses whether to accept the invitation or whether to participate independently in the auction. Agents who join a bidding club first conduct a “pre-auction” within the club; depending on the outcome of the pre-auction some subset of the members of the club bid in the primary auction in a prescribed way. We model this setting as a Bayesian game, including agents’ choices of whether or not to accept a bidding club’s invitation. After describing this general setting, we examine the specific case of bidding clubs for first-price auctions. We show the existence of a Bayes-Nash equilibrium where agents choose to participate in bidding clubs when invited and truthfully declare their valuations to the coordinator. Furthermore, we show that the existence of bidding clubs benefits all agents (including both agent inside and outside of a bidding club) in several different senses.<sup>1</sup>

### 1. INTRODUCTION

The advent of internet markets has spurred new interest in auctions. Most work in both economics and computer science has concentrated on the design of auction protocols from the seller’s perspective, and in particular on optimal (i.e., revenue maximizing) auction design. In this paper we present a class of systems to assist sets of bidders, *bidding clubs*. The idea is similar to the idea behind “buyer clubs” on the Internet (e.g., [www.mobshop.com](http://www.mobshop.com)): to aggregate the market power of individual bidders. Buyer clubs work when buyers’ interests are perfectly aligned; the more buyers join in a purchase the lower the price for everyone. In auctions held on the internet it is relatively easy for multiple agents to cooperate, hiding behind a single auction participant. Intuitively, these bidders can gain by causing others to lower their bids in the case of a first-price auction or by possibly removing the second-highest bidder in the case of a second-price

---

<sup>1</sup>This work was partly supported by DARPA grant number F30602-98-C-0214-P00005.

auction. However, the situation in auctions is not as simple as in buyer clubs, because while bidders can gain by sharing information, the competitive nature of auctions means that bidders’ interests are not aligned. Thus there is a complex strategic relationship among bidders in a bidding club, and bidding club rules must be designed accordingly.

### 1.1. Related Work

While there is relative scarcity of previous work on bidder-centric mechanisms, certainly our work has not been carried out in a vacuum. Below we discuss the most relevant previous work and its relation to ours. This work all comes under the umbrella of *collusion* in auctions, a negative term still reflecting a seller-oriented perspective. We adopt a more neutral stance towards such bidder activities and thus use the term *bidding clubs* rather than the terms *bidding rings* and *cartels* that have been used in the past. However, the technical development is not impacted by such subtle differences in moral attitude.

#### 1.1.1. Collusion in Second-Price Auctions

One of the first formal papers to consider collusion in second-price auctions was written by Graham and Marshall [Graham and Marshall, 1987]. This paper introduces a knockout procedure: agents announce their bids in a pre-auction; only the highest bidder goes to the auction but this bidder must pay a “ring center” the amount of his gain relative to the case where there was no collusion. The ring center pays each agent in advance; the amount of this payment is calculated so that the ring center will budget-balance *ex-ante*, before knowing the agents’ valuations.

Graham and Marshall’s work has been extended to deal with variations in the knockout procedure, differential payments, and relations to the Shapley value [Graham et al., 1990]. The case where only some of the agents are part of the cartel is discussed by Mailath and Zemsky [Mailath and Zemsky, 1991]. Ungern and Sternberg [von Ungern-Sternberg, 1988] discuss collusion in second-price auctions where the designated winner of a cartel is not the agent with the highest valuation. Finally, although this fact is not presented in any existing work of which we are aware, it is also easy to extend Graham and Marshall’s protocol to handle an environment where multiple cartels may operate in the same auction alongside independent bidders.

Overall, a much richer body of work deals with second-price auctions than with first-price auctions. This is possibly explained by the fact that since second-price auctions give rise to dominant strategies, it is possible to study collusion in many settings related to these auctions without performing strategic equilibrium analysis.

### *1.1.2. Collusion in First-Price Auctions*

The key exception to the scarcity of formal work on first-price auctions is a very influential paper by McAfee and McMillan [McAfee and McMillan, 1992]. It is the closest in the literature to our work, and indeed we have borrowed some modelling elements from it. Several sections of their paper, including the discussion of enforcement and the argument for independent private values as a model of agents' valuations, are directly applicable to our paper. However, the setting introduced in their work assumes that a fixed number of agents participate in the auction and that all agents are part of a single cartel that coordinates its behavior in the auction. The authors show optimal collusion protocols for "weak" cartels (in which transfers between agents are not permitted: all bidders bid the reserve price, using the auctioneer's tie-breaking rule to randomly select a winner) and for "strong" cartels (the cartel holds a pre-auction, the winner of which bids the reserve price in the main auction while all other bidders sit out; the winner distributes some of his gains to other cartel members through side payments). A small part of the paper deals with the case where in addition to the single cartel there are also additional agents. However, results are shown only for two cases: (1) when non-cartel members bid without taking the existence of a cartel into account and (2) when each agent  $i$  has valuation  $v_i \in \{0, 1\}$ . The authors explain that they do not attempt to deal with general strategic behavior in the case where the cartel consists of only a subset of the agents; furthermore, they do not consider the case where multiple cartels can operate in the same auction. Finally, a brief presentation of "cartel-formation games" is related to our discussion of agents' decision of whether or not to accept an invitation to join a bidding club.

### *1.1.3. Other Work on Collusion*

Less formal discussion of collusion in auctions can be found in a wide variety of papers. For example, a survey paper that discusses mechanisms that are likely to facilitate collusion in auctions, as well as methods for the detection of such schemes, can be found in [Hendricks and Porter, 1989]. A discussion and comparison of the stability of rings associated with classical auctions can be found in [Robinson, 1985]. That paper concentrates on the case where the valuations of agents in the cartel are honestly reported.

Collusion is also discussed in other settings. For example, the literature discusses collusion that aims to influence purchaser behavior in a repeated procurement setting (see [Feinstein et al., 1985]), and in the context of general Bertrand or Cournot competition (see [Cramton and Palfrey, 1990]).

We should also mention that in an earlier paper we have anticipated some of the results reported here. Specifically, in [Leyton-Brown et al., 2000] we considered bidding clubs under the assumptions that only a single bidding club exists, and that bidders who were not invited to join the club are

not aware of the possibility that a bidding club might exist. The current paper is an extension and generalization of that earlier work.

### 1.2. Distinguishing Features of our Model

Our goal in this work is to study cooperation between self-interested bidders in a rich model that captures many of the characteristics of auctions on the internet. This leads to many differences between our model and models proposed in the work surveyed above (particularly [Graham et al., 1990] and [McAfee and McMillan, 1992]). In particular, we argue that a model of an internet auction setting that includes bidding clubs should include the following features:

1. The number of bidders is stochastic.
2. There is no minimum number of bidders in a bidding club (i.e., bidding clubs are not required to contain all bidders).<sup>2</sup>
3. There is no limit to the number of bidding clubs in a single auction.
4. Club members and independent bidders behave strategically, acting according to correct beliefs about this complex environment.

The first feature above is crucial. In many real-world internet auctions, bidders are not aware of the number of other agents in the economic environment. A bidding club that drops one or more interested bidders is thus undetectable to other bidders in an internet auction. An economic environment with a fixed number of bidders would not model this uncertainty, as the number of interested bidders would be common knowledge among all bidders regardless of the number of bids received in the auction. For this reason, we consider economic environments where the number of bidders is chosen at random. We make use of a model of auctions with stochastic numbers of participants which is due to McAfee and McMillan [McAfee and McMillan, 1987]; we also refer to equilibrium analysis of this model by Harstad, Kagel and Levin [Harstad et al., 1990].

### 1.3. Bidding Clubs at a Glance

Roughly speaking, a scenario with bidding clubs has the following structure:

1. Given a primary auction;
2. Given a set of bidders in that auction, drawn randomly from a set of potential bidders;

---

<sup>2</sup>For technical reasons we will have to assume that there is a finite *maximum* number of bidders in each bidding club; however, this maximum may be any integer greater than or equal to two.

3. Given a partition of bidders into disjoint *clubs*, each of which can be the redundant singleton club;
4. Each bidder chooses whether to bid in the primary auction directly or through his club (it is assumed that this choice is strictly enforceable). In the latter case, the bidder declares his valuation to the club coordinator;
5. Based on the bidders' choices and declarations each club bids in the primary auction, as do both the bidders who elected not to join their respective clubs and the singleton bidders.
6. Each (non-singleton) club bids according to pre-specified, commonly known rules. These rules also specify internal allocations and possible monetary transfers among club members upon the conclusion of the primary auction.

To make bidding clubs a more realistic model of collusion in internet auctions, we restrict bidding club protocols in the following ways:

1. Participation in bidding clubs requires an invitation, but bidders must be free to decline this invitation without (direct) penalty. In this way we include the choice to collude as one of agents' strategic decisions, rather than starting from the assumption that agents will collude.
2. Bidding club coordinators must make money on expectation, and must never lose money. This ensures that third-parties have incentive to run bidding club coordinators. Note that this requirement is not satisfied by a [Graham et al., 1990]-type result, in which bidding clubs (or, in their parlance, cartels) are budget balanced *ex ante*, but may lose money in individual auctions.
3. The bidding club protocol must give rise to an equilibrium where all invited agents choose to participate, even when the bidding club operates in a single auction as opposed to a sequence of auctions. This means that agents can not be induced to collude in a given auction by the threat of being denied future opportunities to collude.

#### 1.4. Overview

This paper consists of two parts. First, sections 2 through 4 present relevant background that does not directly concern cooperation between bidders. In section 2 we give a formal model of an auction with a stochastic number of participants based on the model in [McAfee and McMillan, 1987]. We set up an economic environment in which a finite number of agents is chosen at random from an infinite set of potential agents. We also give a general model of auction mechanisms based on [Monderer and Tennenholtz, 2000],

and define symmetric Bayes-Nash equilibria for the resulting Bayesian game. In section 3 we consider different variations on the first-price auction mechanism. We begin with classical first-price auctions, in which the number of bidders is common knowledge, and then consider first-price auctions in the economic environment from section 2, where the number of bidders is drawn from a known distribution. Combining results from both auction types, we present first-price auctions with participation revelation: auctions in which the number of bidders is stochastic, but the auctioneer announces the number of participants before taking bids. This is the auction mechanism upon which we will base our bidding club protocol for first-price auctions. Finally, section 4 makes use of the revelation principle to show a class of auction mechanisms in which bidders are subject to different payment rules and may have different private information (in addition to their valuations), yet all bid truthfully. We think that this result is interesting in its own right, and certainly it is applicable to settings other than collusion; however, it is also necessary to the proof of the main theorem in section 6.

The second part of our paper is concerned explicitly with bidding clubs, using material from the first part to present a general model of bidding clubs and then a bidding club protocol for first-price auctions. First, section 5 expands the economic environment from section 2 to include the following novel features:

- A finite set of bidding clubs is selected from an infinite set of potential bidding clubs.
- A finite set of agents is selected to participate in the auction, from an infinite set of potential agents. Some agents are associated with bidding clubs, and the whole procedure is carried out in such a way that no agent can gain information about the total number of agents in the economic environment from the fact of his own selection.
- The space of agent types is expanded to include both an agent's valuation, and the number of agents present in that agent's bidding club (equal to one if the agent does not belong to a bidding club).

We introduce notation to describe each agent's beliefs about the number of agents in the economic environment, conditioned on that agent's private information. We also augment the auction mechanism from section 2 to describe additional strategic choices available to agents invited to bidding clubs. In section 6 we examine bidding club protocols for first-price auctions. We begin with two assumptions on the distribution of agent valuations: the first related to continuity of the distribution, and the second to monotonicity of equilibrium bids. After a technical lemma relating equilibrium bids in auctions with stochastic numbers of participants under different distributions, we give a bidding club protocol for first-price auctions with participation revelation. Our main technical results follow:

- We show that it is an equilibrium for agents to accept invitations to join bidding clubs when invited and to disclose their true valuations to their bidding club’s coordinator. Under the same equilibrium, singleton agents bid as they would in an auction with a stochastic number of participants in an economic environment without bidding clubs, in which the distribution over the number of participants is the same as in the bidding clubs setting.
- In equilibrium each agent is better off as a result of his own club (that is, his expected payoff is higher than would have been the case if his club never existed, but other clubs—if any—still did exist).
- In equilibrium each club increases all non-members’ expected payoffs, as compared to equilibrium in the case where all club members participated in the auction as singleton bidders, but all other clubs—if any—still existed.
- In equilibrium each agent’s expected payoff is identical to the case in which no clubs exist; note that since clubs make money on expectation, if clubs are willing to make money (or break even) only on expectation, they could distribute some of their *ex ante* expected profits among the club members, ensuring that all bidders gain on expectation.

Finally, sections 7 and 8 consist of discussion and conclusions. We touch on questions of trustworthiness of coordinators, legality of bidding clubs and steps an auctioneer could take to disrupt the operation of bidding clubs in her auction.

## 2. AUCTION MODEL

In this section we provide a (non-controversial) auction model, meant to capture an internet auction setting such as eBay. Of course, this model is applicable to many other auctions as well. Auctions may be seen as consisting of an economic environment plus an auction mechanism which together define a Bayesian game. First, our economic environment consists of a stochastic number of agents, each of which has private information about the number of participants in the auction and knows the distribution from which others’ types are drawn. This section draws heavily on work by McAfee and McMillan [McAfee and McMillan, 1987] on auctions with a stochastic number of participants. Second, the game includes an auction mechanism in which the agents participate; this section is based on [Monderer and Tennenholtz, 2000]. After defining these elements, we give a formal definition of the Bayesian game.

## 2.1. The Economic Environment

An economic environment  $E$  consists of a finite set of agents who have non-negative valuations for a good at auction, and a distinguished agent 0—the seller or center. The set of agents is selected by an exogenous process, and each agent is unaware of the total number of agents participating in the economic environment. Following [McAfee and McMillan, 1987], let the set of agents who may participate in the economic environment be  $\mathcal{A} \equiv \mathbb{N}$ . Let  $\beta_A$  represent the probability that a finite set  $A \subset \mathcal{A}$  is the set of agents. The probability that  $n$  agents<sup>3</sup> will participate in the auction is  $\gamma_A(n) = \sum_{A, |A|=n} \beta_A$ . All agents know the probability distribution  $\beta_A$ . Once an agent  $k$  is selected, he updates his probability of the number of agents present as:

$$p_n^k = \frac{\sum_{A, |A|=n, k \in A} \beta_A}{\sum_{A, k \in A} \beta_A}. \quad (1)$$

We deviate from the model in [McAfee and McMillan, 1987] by adding the assumption that it is common knowledge that all bidders are equally likely to be chosen. Hence  $p_n^k$  is the same for all  $k$ ; we will hereafter refer only to  $p_n$ . Finally, we assume that  $\gamma_A(0) = \gamma_A(1) = 0$ ; at least two agents will participate in the auction.

Let  $\mathcal{T}$  be the set of possible agent types. The type  $\tau_i \in \mathcal{T}$  of agent  $i$  is the tuple  $(v_i, s_i) \in V \times \mathcal{S}$ .  $v_i$  denotes an agent's valuation: his maximal willingness to pay for the good offered by the center. We assume that  $v_i$  represents a purely private valuation for the good, and that  $v_i$  is selected independently from the other  $v_j$ 's of other agents from a known distribution,  $F$ , having density function  $f$ . By  $s_i$  we denote agent  $i$ 's signal: his private information about the number of agents in the auction. In this section we will consider the simple case where  $\mathcal{S} = \{\emptyset\}$ : it is common knowledge that all agents receive the null signal, and hence gain no additional information about the number of agents. Note, however, that the economic environment itself is always common knowledge, and so agents always have some information about the number of agents even when they receive the null signal. We will consider more complex signals in section 5. We will use the notation  $p_n^{\tau_i}$  to denote the probability that agent  $i$  assigns to there being  $n$  agents in the auction, conditioned on his type  $\tau_i$ . Throughout the paper we will use uppercase  $P$  to denote the whole probability distribution as compared to the probability of a particular number of agents which we have denoted by lowercase  $p$ ; in this case we denote the whole distribution conditioned on  $i$ 's type as  $P^{\tau_i}$ .

The utility function of agent  $i$ ,  $u_i : \mathbb{R} \rightarrow \mathbb{R}$  is linear, normalized with  $u_i(0) = 0$ . The utility of agent  $i$  (having valuation  $v_i$ ) when asked to pay

---

<sup>3</sup>When we say that  $n$  agents participate in the auction we do not count the distinguished agent 0, who is always present.

$t$  is  $v_i - t$  if  $i$  is allocated a good, and it is 0 otherwise. Thus, we assume that there are no externalities in agents' valuations and that agents are risk-neutral.

## 2.2. The Auction Mechanism

We denote the possible allocations of the good to the agents by  $\Pi$ . An auction mechanism is a tuple  $(\mathcal{M}, g, t)$ , where:

- $\mathcal{M}$  is the set of possible messages an agent may send.
- $g : \mathcal{M}^n \rightarrow \Delta(\Pi)$  is an allocation function where  $\Delta(\Pi)$  is the tuple of distribution functions over  $\Pi$  (e.g., the allocation may include random elements).
- $t = (t_1, t_2, \dots, t_n)$ ,  $t_i : \mathcal{M}^n \times \Pi \rightarrow \mathbb{R}$  is the (monetary) transfer function for agent  $i$ .

Notice that  $n$  is a parameter. Technically, an auction mechanism defines  $g$  and  $t$  for any number of participants, and can be therefore considered as a set of tuples (one for each number of agents).

Given the above, the dynamics of an auction mechanism can be described as follows:

- Each agent  $i$  sends a message  $\mu_i$  to the center. We denote the set of messages received by the center as  $\mu$ .
- The center conducts a lottery according to the distribution  $g(\mu)$ , and selects the allocation  $\pi$ .
- Agent  $i$  gets  $\pi_i$ , and is required to transfer  $t_i(\mu, \pi)$  to the center.
- The utility of  $i$  is  $v_i - t_i(\mu, \pi)$  if he is assigned a good, and it is  $-t_i(\mu, \pi)$  otherwise.

## 2.3. The Bayesian Game

The auction mechanism  $(\mathcal{M}, g, t)$ , in conjunction with the economic environment  $E$ , defines a Bayesian game. We will use the following definitions and notation. A strategy  $b_i : \mathcal{T} \rightarrow \mathcal{M}$  for agent  $i$  is a mapping from his type  $\tau_i$  to a message  $\mu_i$ . This may be the null message, which means that he has elected not to participate in the auction.  $\Sigma$  denotes the set of possible strategies, i.e., the set of functions from types to messages in  $\mathcal{M}$ . Each agent's type is that agent's private information, but the whole setting is common knowledge.

For notational simplicity we only define symmetric equilibria, where all agents bid the same function of their type, as this is sufficient for our purposes in this paper. A more general definition would proceed along the

same lines. By  $L_i(\tau_i, b_i, b^{j-1})$  we denote agent  $i$ 's *ex post* expected utility given that his type is  $\tau_i$ , he follows the strategy  $b_i$  and all other agents use the strategy  $b$ , in the case that there are a total of  $j$  agents. The strategy profile  $b^n \in \Sigma^n$  is a symmetric equilibrium if and only if:

$$\forall i \in A, \forall \tau_i \in \mathcal{T}, b \in \operatorname{argmax}_{b_i \in \Sigma} \sum_{j=2}^{\infty} p_j^{\tau_i} L_i(\tau_i, b_i, b^{j-1}) \quad (2)$$

### 3. FIRST-PRICE AUCTIONS

In this section we discuss several different variants of the first-price auction. First we describe classical first-price auctions, in which a fixed number of participants belong to the economic environment, and hence the number of bidders is common knowledge. Next we consider first-price auctions with a stochastic number of participants, where the number of bidders in the economic environment is drawn from a known distribution. Using the previous two settings, we present first-price auctions with participation revelation, where the number of agents is chosen stochastically, but the auctioneer announces the number of agents who have registered in the auction before taking bids. This last type of first-price auction is the one we will consider in our discussion of bidding clubs in section 6.

#### 3.1. Classical first-price auctions

In a classical first-price auction, each participant submits a bid in a sealed envelope. The agent with the highest bid wins the good and pays the amount of his bid, and all other participants pay nothing. In the case of a tie, the winner of the auction is selected uniformly at random from the bidders who tied for the highest bid. (Note, however, that when  $F$  is continuous and has no atoms the probability of two bidders having the same type is 0; ties will therefore occur with probability 0 if bidders follow an equilibrium in which they all bid a strictly monotonically-increasing function of their valuations.) The equilibrium analysis of first-price auctions is quite standard:

**PROPOSITION 1.** *If valuations are selected independently according to the uniform distribution on  $[0, 1]$  then it is a symmetric equilibrium for each agent  $i$  to follow the strategy:*

$$b(v_i) = \frac{n-1}{n} v_i.$$

Using classical equilibrium analysis (e.g., following Riley and Samuelson [Riley and Samuelson, 1981]) it is possible to show how classical first-price auctions can be generalized to an arbitrary continuous distribution  $F$ .

**PROPOSITION 2.** *If valuations are selected from a continuous distribution  $F$  then it is a symmetric equilibrium for each agent  $i$  to follow the strategy:*

$$b(v_i) = v_i - F(v_i)^{-(n-1)} \int_0^{v_i} F(u)^{n-1} du.$$

In both cases, observe that although  $n$  is a free variable,  $n$  is not a parameter of the strategy; the same is true of the distribution  $F$ . Agents deduce this information from their full knowledge of the economic environment. It is useful, however, to have notation specifying the amount of the equilibrium bid as a function of both  $v$  and  $n$ . We write

$$b^e(v_i, n) = v_i - F(v_i)^{-(n-1)} \int_0^{v_i} F(u)^{n-1} du. \quad (3)$$

### 3.2. First-price auctions with a stochastic number of bidders

In the economic environment described in section 2.1 the number of agents is not a constant; rather, it is chosen stochastically from a known probability distribution. An equilibrium for this setting was demonstrated by Harstad, Kagel and Levin [Harstad et al., 1990]:

**PROPOSITION 3.** *If valuations are selected from a continuous distribution  $F$  and the number of bidders is selected from the distribution  $P$  then it is a symmetric equilibrium for each agent  $i$  to follow the strategy:*

$$b(v_i) = \sum_{j=2}^{\infty} p_j b^e(v_i, j)$$

Observe that  $b^e(v_i, j)$  is the amount of the equilibrium bid for a bidder with valuation  $v_i$  in a setting with  $j$  bidders as described in section 3.1 above.  $P$  is deduced from the economic environment.<sup>4</sup> We overload our previous notation for the equilibrium bid, this time as a function of the agent's valuation and the probability distribution  $P$ . Thus we write:

$$b^e(v_i, P) = \sum_{j=2}^{\infty} p_j b^e(v_i, j) \quad (4)$$

We will make frequent use of this function throughout the paper. An important note is that it describes the equilibrium bid in the situation where the economic environment is such that the number of agents is chosen by  $P$  and where all agents receive the null signal.

---

<sup>4</sup>Recall that  $P$  is a set:  $p_j \in P$  for all  $j \geq 0$ , where  $p_j$  denotes the probability that the economic environment contains exactly  $j$  agents.

### 3.3. First-price auctions with participation revelation

In some first-price auctions (e.g., auctions held on the internet), bidders participate in an economic environment where the number of bidders in the auction is not common knowledge. However, this can be helpful information for bidders. One obvious way of addressing this problem is to introduce a two-phase mechanism with revelation of the number of participants between the stages. Specifically, a first-price auction with participation revelation is as follows:

1. Agents indicate their intention to bid in the auction.
2. The auctioneer announces  $n$ , the number of agents who registered in the first phase.
3. Agents submit bids to the auctioneer. The auctioneer will only accept bids from agents who registered in the first phase.
4. The agent who submitted the highest bid is awarded the good for the amount of his bid; all other agents are made to pay 0.

It is unsurprising that, although a first-price auction with participation revelation may have a stochastic number of participants,

*PROPOSITION 4. There exists an equilibrium of the first-price auction with participation revelation where every agent  $i$  indicates the intention to participate and bids according to  $b^e(v_i, n)$ .*

*Proof.* Agents are always better off participating in first-price auctions as long as there is no participation fee. The only way of participating is to declare the intention to participate in the first phase of the auction. Thus the number of agents announced by the auctioneer is equal to the total number of agents in the economic environment. From proposition 2 it is best for agent  $i$  to bid  $b^e(v_i, n)$  when it is common knowledge that the number of agents in the economic environment is  $n$ . That is exactly the case under our mechanism. ■

In section 6 we will be concerned with first-price auctions with information revelation, but we will show an equilibrium in which the number of agents registering in the first phase is smaller than the total number of agents participating in the auction, because some bidders with low valuations drop out as part of a collusive agreement. The auctioneer's declaration acts as a signal about the total number of bidders, but individual agents will still be uncertain about the total number of opponents they face.

## 4. TRUTHFUL EQUILIBRIA IN ASYMMETRIC MECHANISMS

In this section we describe a particular class of auction mechanisms that are asymmetric in the sense that every agent is subject to the same

allocation rule but to a potentially different payment rule, and furthermore that agents may receive different signals. It will be helpful for the proof of our main theorem in section 6 to show that a truth-revealing equilibrium exists in such auctions under the following two conditions:

1. The auction allocates the good to the agent who submits the highest bid.
2. Consider the auction  $M_i$  in which *all* agents are subject to agent  $i$ 's payment rule and the above allocation rule, and where (hypothetically) *all* agents receive the signal  $s_i$ .<sup>5</sup> Truth-revelation is a symmetric equilibrium in  $M_i$ .

Observe that the second condition above is less restrictive than it may appear. From the revelation principle we can see that for every auction with a symmetric equilibrium there is a corresponding auction in which truth-revealing is an equilibrium that gives rise to the same allocation and the same payments for all agents.  $M_i$  can thus be seen as a revelation mechanism for some other auction that has a symmetric equilibrium.

More formally, given a good  $g$ , let  $\mathbf{M}$  represent a set of auctions  $\{M_1, \dots, M_n\}$  which all allocate the good to the agent who submits the highest bid, and which are all truth-revealing direct mechanisms for  $n$  risk-neutral agents with independent private valuations drawn from the same distribution. We now define another auction  $\bar{M}$ :

1. Each agent  $i$  sends a message  $\mu_i$  to the center.
2. The center allocates the good to the agent  $i$  with  $\mu_i \in \max_j \mu_j$ . If multiple agents submit the highest message, the tie is broken in some arbitrary way.
3. Agent  $i$  is made to transfer  $t_i(\mu, \pi)$  to the center.<sup>6</sup> The transfer function  $t_i$  is taken from  $M_i \in \mathbf{M}$ .

We can now show:

LEMMA 1. *Truth-revelation is an equilibrium of  $\bar{M}$ .*

*Proof.* The payoff of agent  $i$  is uniquely determined by the allocation rule, the transfer function  $t_i$ , and all agents' strategies. Assume that the other agents are truth revealing, then the other agents' behavior, the allocation rule, and agent  $i$ 's payment rule are all identical in  $\bar{M}$  and  $M_i$ . Since truth-revelation is an equilibrium in  $M_i$ , truth-revelation is agent  $i$ 's best response in  $\bar{M}$ . ■

---

<sup>5</sup>That is, for every agent  $j$  in the real auction, we create an agent  $k$  in the hypothetical auction  $M_i$  having type  $\tau_k = (v_j, s_i)$ .

<sup>6</sup>Of course, this transfer can be either positive or negative.

*Example.* Consider an auction for a single good  $g$ , where eight agents bid for the good. The agents' valuations are IPV, IID from a known distribution  $F$ , and the agents are risk-averse. Let  $M_1$  be a revelation mechanism for a first-price auction: i.e., agents declare their valuations, and the winner is charged  $b^e(v, 8)$ . In an economic environment consisting of eight agents with IPV valuations from  $F$  it is an equilibrium of  $M_1$  for agents to truthfully declare their valuations to the center. Let  $M_2$  be a second-price auction; truthful declaration is a weakly dominant strategy under this auction type. Both  $M_1$  and  $M_2$  allocate the good to the agent with the highest declaration, and so these auctions meet the conditions given at the beginning of the section. Now consider an auction  $\bar{M}$  where odd-numbered agents are subject to the payment rule from  $M_1$ , and even-numbered agents are subject to the payment rule from  $M_2$ . By lemma 1, truth-revelation is an equilibrium of  $\bar{M}$ . There are other differences between payment rules that can cause agents' expected utilities to differ: for example, lemma 1 would still hold if  $M_2$  gave each agent an additional payment of \$10 for participating in the auction.

The next corollary, which follows directly from the lemma, compares a single agent's expected utility under two different auctions  $\bar{M}$  and  $\bar{M}'$ , which implement different payment rules. We will need this result for our proof of theorem 1.

**COROLLARY 1.** *Consider two auctions  $\bar{M}$  and  $\bar{M}'$ , defined as above, which both implement the same transfer function for agent  $i$ . Agent  $i$ 's expected utility is the same in both  $\bar{M}$  and  $\bar{M}'$ .*

*Proof.* The payoff of agent  $i$  is uniquely determined by the allocation rule, its transfer function, and all agents' strategies. Both  $\bar{M}$  and  $\bar{M}'$  have the same allocation rule. Lemma 1 tells us that truth revelation is a best response for all agents in both  $\bar{M}$  and  $\bar{M}'$ , so all agents' strategies are identical in the two auctions. In general, agents may not receive the same expected utility from  $\bar{M}$  and  $\bar{M}'$ . However, since  $i$  has the same transfer function in both auctions,  $i$ 's expected utility in  $\bar{M}$  is equal to his expected utility in  $\bar{M}'$ . ■

## 5. AUCTION MODEL FOR BIDDING CLUBS

In this section we extend both the economic environment and auction mechanism from section 2 to include the characteristics necessary for a model of bidding clubs. Because our aim is not to model a situation where agents' *decision* to collude is exogenous—as this would gloss over the question of whether the collusion is stable—we include the collusive protocol as part of the model and show that it is individually rational *ex post* (i.e., after agents have observed their valuations) for agents to choose to collude. However, we do consider exogenous the selection of the set of agents who are *offered* the opportunity to collude. Furthermore, we want to show the

impact of the possibility of collusion upon non-colluding agents; indeed, even colluding agents must take into account the possibility that *other* groups of agents in the auction may also be colluding. Once we have defined the new economic environment and auction mechanism, a well-defined Bayesian game will be specified by every tuple of primary auction type, bidding club rules and distributions of agent types, the number of agents and the number of bidding clubs.

### 5.1. The Economic Environment

We extend the economic environment  $E$  from the previous section to consist of a set of agents who have non-negative valuations for a good at auction, the distinguished agent 0 and a set of bidding club coordinators who may invite agents to participate in a bidding club. Intuitively, we construct an environment where an agent's belief update after observing the number of agents in his bidding club does not result in any change in the distribution over the number of *other* agents in the auction, because the number of agents in each bidding club is independent of the number of agents in every other bidding club.

#### 5.1.1. Coordinators

Coordinators are not free to choose their own strategies; rather, they act as part of the mechanism for a subset of the agents in the economic environment. We select coordinators in a process analogous to our previous approach for exogenously selecting agents: we draw a finite set of individuals from an infinite set of potential coordinators. In this case, however, this finite set is considered “potential coordinators”; in section 5.1.2 we will describe which potential coordinators are “actualized”, i.e., correspond to actual coordinators. Possible coordinators that are not actualized will correspond to singleton bidders in the auction.

More formally, let  $\mathcal{C} \equiv \mathbb{N}$  (excluding 0) be the set of all coordinators.  $\beta_C$  represents the probability that a finite set  $C \subset \mathcal{C}$  is selected to be the set of potential coordinators. We add the restriction that all coordinators are equally likely to be chosen. A consequence of this restriction is that an agent's knowledge of the coordinator with whom he is associated does not give him additional information about what other coordinators may have been selected. We denote the probability that an auction will involve  $n_c$  potential coordinators as  $\gamma_C(n_c) = \sum_{C, |C|=n_c} \beta_C$ . The distribution  $\beta_C$  is common knowledge. We assume that  $\gamma_C(0) = \gamma_C(1) = 0$ : at least two potential coordinators will be associated with each auction.

#### 5.1.2. Agents

We independently associate a random number of agents with each potential coordinator, again drawing a finite set of actual agents from an

infinite set of potential agents. If only one (actual) agent is associated with a potential coordinator, the potential coordinator will not be actualized and hence the agent will not belong to a bidding club. In this way we model agents who participate directly in the auction without being associated with a coordinator. If more than one agent is associated with a potential coordinator, the coordinator *is* actualized and all the agents receive an invitation to participate in the bidding club.

More formally, let  $\mathcal{A} \equiv \mathbb{N}$  be the set of all agents, and let  $\kappa$  be the maximum number of agents who may be associated with a single bidding club. Partition  $\mathcal{A}$  into subsets, where agent  $i$  belongs to the subset  $\mathcal{A}_{\lceil i/\kappa \rceil}$ . Let  $\beta_A$  be the probability that a finite set  $A \subset \mathcal{A}_i$  is the set of agents associated with potential coordinator  $i$ ; we assume that this distribution is the same for all  $i$ . Furthermore, as above, we assume that it is common knowledge that all agents are equally likely to be chosen. The probability that  $n$  agents will be associated with a potential coordinator is denoted  $\gamma_A(n) = \sum_{A, |A|=n} \beta_A$ . By the definition of  $\kappa$ ,  $\forall j > \kappa, \gamma_A(j) = 0$ ; we assume that  $\gamma_A(0) = 0$  and that  $\gamma_A(1) < 1$ .

#### 5.1.3. Signals

Each agent receives a signal informing him of the number of agents in his bidding club; as above we denote this signal as  $s_i$ .<sup>7</sup> Of course, if this number is 1 then there is no coordinator for the agent to deal with, and he will simply participate in the main auction. Note also that agents are neither aware of the number of potential coordinators for their auction nor the number of actualized potential coordinators, though they are aware of both distributions.

#### 5.1.4. Beliefs

Once an agent is selected, he updates his probability distribution over the number of actual agents in the economic environment. Not all agents will have the same beliefs—agents who have been signaled that they belong to a bidding club will expect a larger number of agents than singleton agents. We denote by  $p_m^{n,k}$  the probability that there are a total of  $m$  agents in the auction, given that there are  $n$  bidding clubs and that there are  $k$  agents in the bidder's own club; we denote the whole distribution  $P^{n,k}$ . Because the numbers of agents in each bidding club are independent, observe that every agent in the whole auction has the same beliefs about the number of other agents in the economic environment, discounting those agents in his own bidding club. Hence agent  $i$ 's beliefs are described by

---

<sup>7</sup>In fact, none of our results require that agents know the *number* of agents in their bidding clubs; it would be sufficient that agents know *whether* they belong to a bidding club. We consider the setting where agents' signals are more informative because it simplifies the exposition of the main theorem.

the distribution  $P^{n,s_i}$ . It is important to note that  $P^{n,s_i}$  is simply another distribution over the number of agents in the auction. Although this shorthand makes reference to the bidding club economic environment in order to describe the construction of the distribution, it makes sense to talk about a classical auction with a stochastic number of bidders (i.e., section 3.2) where the number of bidders is distributed according to  $P^{n,k}$  for given values of  $n$  and  $k$ .

### 5.2. The Augmented Auction Mechanism

Bidding clubs, in combination with a main auction, induce an augmented auction mechanism for their members:

1. A set  $A$  of bidders is invited to join the bidding club.
2. Each agent  $i$  sends a message  $\mu_i$  to the bidding club coordinator. This may be the null message, which indicates that the agent will not participate in the coordination and will instead participate freely in the main auction. Otherwise, agent  $i$  agrees to be bound by the bidding club rules, and  $\mu_i$  is agent  $i$ 's declared valuation for the good. Of course,  $i$  can lie about his valuation.
3. Based on pre-specified and commonly-known rules, and on the information all the members supply, the coordinator selects a subset of the agents to bid in the main auction. The coordinator may bid on behalf of these agents (e.g., using their ID's on the auction web site) or it may instruct agents on how to bid. In either case we assume that the coordinator can force agents to bid as desired, for example by imposing a charge on agents who do not behave as directed.
4. If a bidder represented by the coordinator wins the main auction, he is made to pay the amount required by the auction mechanism to the auctioneer. In addition, he may be required to make an additional payment to the coordinator.

Any number of coordinators may participate in an auction. However, we assume that there is only a single coordination protocol, and that this protocol is common knowledge.

## 6. BIDDING CLUBS FOR FIRST-PRICE AUCTIONS

In this section we first give some (mild) assumptions about the distribution of agent valuations, then use these assumptions to prove a technical lemma. We then give the bidding club protocol for first-price auctions. We consider a first-price auction with participation revelation as described in section 3.3. Bidders indicate their intention to participate, the auctioneer announces the total number of bidders and then bidders place their

bids. The bidding club decides whether to drop bidders before the first phase; therefore the number announced by the auctioneer does not include dropped bidders. We show an equilibrium of this auction, and demonstrate that agents gain under this equilibrium.

### 6.1. Assumptions

Our results hold for a broad class of distributions of agent valuations—all distributions for which the following two assumptions are true.

First, we assume that  $F$  is continuous and atomless.

In order to give our second assumption, we must introduce some notation. Define:

$$P_{x \geq i} = \sum_{x=i}^{\infty} p_x. \quad (5)$$

We now define the relation “ $<$ ” for probability distributions:

$$P < P' \text{ iff } \exists l (\forall i < l, P_{x \geq i} = p'_{x \geq i} \text{ and } \forall i \geq l, P_{x \geq i} < P'_{x \geq i}). \quad (6)$$

We are now able to state our second assumption:

$$(P < P') \text{ implies that } \forall v, b^e(v, P) < b^e(v, P'), \quad (7)$$

Intuitively, we assume that every agent’s symmetric equilibrium bid in a setting with a stochastic number of participants drawn from  $P'$  is strictly greater than that agent’s symmetric equilibrium bid in a setting with a stochastic number of participants drawn from  $P$ , in the case where  $P'$  stochastically dominates  $P$ .

### 6.2. A Technical Lemma

Recall from section 5.1.4 that the notation  $P^{n,k}$  may be seen as defining a probability distribution over the number of agents that is independent of the bidding club setting. It is thus possible to discuss equilibrium bids in the classical stochastic settings where the number of bidders is drawn from such a distribution. While it will remain to show why these values are meaningful in our setting where (among other differences) agents have asymmetric information, it will be useful to prove the following lemma about the classical stochastic setting:

$$\text{LEMMA 2. } \forall k \geq 2, \forall n \geq 2, \forall v, b^e(v, P^{n+k-1,1}) > b^e(v, P^{n,k})$$

*Remark.* For convenience and to preserve intuition in what follows we will refer to the number of potential coordinators and the number of agents belonging to a coordinator even though we concern ourselves with

the classical economic environment from section 2.1 where bidding clubs do not exist. The number of potential coordinators is shorthand for the number  $n_c$  drawn from  $\gamma_C$  in the first phase of the procedural definition of the distribution  $P^{n,k}$ . Likewise the number of agents associated with a potential coordinator is shorthand for the number of agents chosen from one of the  $n_c$  iterative draws from  $\gamma_A$ . Intuitively, this lemma asserts that the symmetric equilibrium bid is always higher when more agents belong to the main auction as singleton bidders and the total number of agents is held constant.

*Proof.* Recall our second assumption from section 6.1. We defined  $P < P'$  as the proposition that  $\exists l (\forall i < l, P_{x \geq i} = P'_{x \geq i} \text{ and } \forall i \geq l, P_{x \geq i} < P'_{x \geq i})$ . Our second assumption was that  $(P < P')$  implies that  $\forall v, b^e(v, P) < b^e(v, P')$ . It is thus sufficient to show that  $P^{n+k-1,1} > P^{n,k}$ . We will take  $l = n + k$ .

First we will show that  $\forall j < n + k, P_{x \geq j}^{n+k-1,1} = P_{x \geq j}^{n,k}$ . The distribution  $P^{n+k-1,1}$  expresses the belief that there are  $n+k-2$  potential coordinators, the membership of which is distributed as described in section 5.1, and one potential coordinator that is known to contain only a single bidder. The distribution  $P^{n,k}$  expresses the belief that there are  $n-1$  potential coordinators, the membership of which is again distributed as described in section 5.1, and one potential coordinator that is known to contain exactly  $k$  bidders. Under both distributions it is certain that there are at least  $n+k-1$  agents. Therefore  $\forall j < n+k, P_{x \geq j}^{n+k-1,1} = P_{x \geq j}^{n,k} = 1$ .

Second,  $\forall j \geq n+k, P_{x \geq j}^{n+k-1,1} > P_{x \geq j}^{n,k}$ . Considering  $P^{n+k-1,1}$ , observe that for  $n+k-2$  of the potential coordinators the probability that this coordinator contains a single agent is less than one and these probabilities are all independent; the last potential coordinator contains a single agent with probability one. Considering  $P^{n,k}$ , there are  $n-1$  potential coordinators where the probability of containing a single agent is less than one, exactly as above, and  $k$  potential coordinators certain to contain exactly one agent. Thus the two distributions agree exactly about  $n-1$  of the potential coordinators, which both hold to contain more than a single agent, and likewise both distributions agree that one of the potential coordinators contains exactly one agent. However, there remain  $k-1$  potential coordinators about which the distributions disagree;  $P^{n+k-1,1}$  always generates a greater or equal number of agents for these potential coordinators, as compared to  $P^{n,k}$ . Under the latter distribution all these agents are singletons with probability one, while under the former there is positive probability that each of the potential coordinators contains more than one agent. As long as  $k \geq 2$ , there is at least one potential coordinator for which  $P^{n+k-1,1}$  stochastically dominates  $P^{n,k}$ . Thus  $\forall k \geq 2, \forall n \geq 2, \forall v P^{n+k-1,1} > P^{n,k}$ . ■

### 6.3. First-Price Auction Bidding Club Protocol

What follows is the protocol of a coordinator who approaches  $k$  agents.

1. Each agent  $i$  sends a message  $\mu_i$  to the coordinator.
2. If at least one agent declines participation then the coordinator registers in the main auction for every agent who accepted the invitation to the bidding club. For each bidder  $i$ , the coordinator submits a bid of  $b^e(\mu_i, P^{n,k})$ , where  $n$  is the number of bidders announced by the auctioneer.
3. If all  $k$  agents accepted the invitation then the coordinator drops all bidders except the bidder with the highest reported valuation, who we will denote as bidder  $h$ . For this bidder the coordinator will place a bid of  $b^e(\mu_h, P^{n,1})$  in the main auction.
4. If bidder  $h$  wins in the main auction, he is made to pay  $b^e(\mu_h, P^{n,1})$  to the center and  $b^e(\mu_h, P^{n,k}) - b^e(\mu_h, P^{n,1})$  to the coordinator.

We are now ready to prove the main theorem of the paper:

**THEOREM 1.** *It is an equilibrium for all bidding club members to choose to participate and to truthfully declare their valuations to their respective bidding club coordinators, and for all non-bidding club members to participate in the main auction with a bid of  $b^e(v, P^{n,1})$ .*

*Proof.* We first prove that the above strategy is in equilibrium for both categories of bidders given that agents all participate; we then prove that participation is rational for all agents.

For the proof of equilibrium we consider a one-stage mechanism which behaves as follows:

1. The center announces  $n$ , the number of bidders in the main auction.
2. Bidders submit bids (messages) to the mechanism.
3. The bidder with the highest bid is allocated the good.
4. The winning bidder is made to pay  $b^e(v_i, P^{n,s_i})$ .

This one-stage mechanism has the same payment rule for bidding club bidders as the bidding club protocol given above, but no longer implements a first-price payment rule for singleton bidders. In order to prove that the strategies given in the statement of the theorem are an equilibrium, it is sufficient to show that truthful bidding is an equilibrium for all bidders under the one-stage mechanism. Observe that this mechanism may be seen as a mechanism  $\bar{M}$  in the sense of lemma 1: it allocates the good to the agent who submits the highest message, and (by definition of  $b^e$ ) the

auction  $M_i$  in which *all* agents are subject to agent  $i$ 's payment rule and receive the signal  $s_i$  has truth revelation as a symmetric equilibrium.

*Strategy of non-club bidder:* Assume that all bidding club agents bid truthfully. Further assume that all non-club agents also bid truthfully except for agent  $i$ . The probability distribution  $P^{n,1}$  correctly describes the beliefs of non-club agents, given the auctioneer's announcement that there are  $n$  bidders in the main auction. Although agents in bidding clubs have additional information about the number of agents—each agent knows that there is at least one other agent in his own club—their prescribed behavior is to place bids of  $b^e(\mu, P^{n,1})$  in the main auction. Agent  $i$  thus faces a stochastic number of agents distributed according to  $P^{n,1}$  and all bidding  $b^e(v, P^{n,1})$ . Using the result from lemma 1,  $i$ 's strategic decision is the same as under a mechanism where all agents are subject to his payment rule and share his signal  $s_i$ , and with a stochastic number of bidders distributed according to  $P^{n,1}$ . In particular, it does not matter that the club members are subject to different payment rules and have additional information, and so  $i$  will also bid  $b^e(v, P^{n,1})$ .

*Strategy of club bidder:* Assume that all agents accept the invitation to join their respective clubs and then truthfully declare their valuations, excluding agent  $i$  who decides to participate but considers his bid. Once again, observe that  $i$  is in a setting that is exactly described by lemma 1:  $P^{n,k}$  really does describe the distribution over the number of agents given his signal, and the bidder submitting the highest (global) message will always be allocated the good. Therefore the information asymmetry does not affect  $i$ 's strategy, and so truthful bidding is a best response for agent  $i$ .

We now turn to the question of participation; for this part of the proof we return to the original, multi-stage mechanism.

*Participation of non-club bidder:* Because there is no participation fee, it is always rational for a bidder to participate in a first-price auction.

*Participation of club bidder:* Likewise, because there is no participation fee, all bidding club bidders will participate in the auction, but must decide whether or not to accept their coordinators' invitations. Assume that all agents except for  $i$  join their respective clubs and bid truthfully, and agent  $i$  must decide whether or not to join his bidding club. Agent  $i$  knows the number of agents in his bidding club and updates his distribution over the number of agents in the whole auction as  $P^{n,k}$ .

Consider the classical stochastic case where all bidders have the same information as  $i$  (and are subject to the same payment rules): from proposition 3 it is a best response for  $i$  to bid  $b^e(v_i, P^{n,k})$ . In this setting  $i$ 's expected gain is the same as in the equilibrium where all bidding club members (including  $i$ ) join their clubs and bid truthfully, by corollary 1.

As a result of  $i$  declining the offer to participate in the bidding club there are  $n - 1$  bidders in the main auction placing bids of  $b^e(v, P^{n+k-1,1})$  and  $k - 1$  other bidders placing bids of  $b^e(v, P^{n,k})$ . Note that this occurs

because the singleton bidders and other bidding clubs in the main auction follow a strategy that depends on the number of bidders announced by the auctioneer; hence they bid as though all the  $k - 1$  bidders from the disbanded bidding club might each be independent bidding clubs. We know from lemma 2 that  $b^e(v, P^{n+k-1,1}) > b^e(v, P^{n,k})$ . Thus the singleton bidders and other bidding clubs will bid a higher function of their valuations than the bidders from the disbanded bidding club. It always reduces a bidder's expected gain in a first-price auction to cause other bidders to bid above the equilibrium, because it reduces the chance that he will win without affecting his payment if he does win. This is exactly the effect of  $i$  declining the offer to join his bidding club: the  $k - 1$  other bidders from  $i$ 's bidding club bid according to the equilibrium of the classical stochastic case discussed above, but the  $n - 1$  singleton and bidding club bidders submit bids that exceed the symmetric equilibrium amount. Therefore  $i$ 's expected gain is smaller if he declines the offer to participate than if he accepts it. ■

#### 6.4. Do bidding clubs cause agents to gain?

We can show that bidders are better off being invited to a bidding club than being sent to the auction as singleton bidders. Intuitively, an agent gains by not having to consider the possibility that other bidders who would otherwise have belonged to his bidding club might themselves be bidding clubs.

**THEOREM 2.** *An agent  $i$  has higher expected utility in a bidding club of size  $k$  bidding as described in theorem 1 than he does if the bidding club does not exist and  $k$  additional agents (including  $i$ ) participate directly in the main auction as singleton bidders, again bidding as described in theorem 1.*

*Proof.* Consider the counterfactual case where agent  $i$ 's bidding club does not exist, and all the members of this bidding club become singleton bidders. We will show that  $i$  is better off as a member of the bidding club than in this case. If there were  $n$  potential coordinators in the original auction and  $k$  agents in  $i$ 's bidding club, then the auctioneer will announce  $n + k - 1$  as the number of participants in the new auction. Under the equilibrium from theorem 1, as a singleton bidder  $i$  will bid  $b^e(v_i, P^{n+k-1,1})$ . If he belonged to the bidding club and followed the same equilibrium  $i$  would bid  $b^e(v_i, P^{n,k})$ . In both cases the auction is economically efficient, which means  $i$  is better off in the auction that requires him to pay a smaller amount when he wins. Lemma 2 shows that  $\forall k \geq 2, \forall n \geq 2, \forall v, b^e(v, P^{n+k-1,1}) > b^e(v, P^{n,k})$ , and so our result follows. ■

We can also show that singleton bidders and members of other bidding clubs benefit from the existence of each bidding club in the same sense. Fol-

lowing an argument similar to the one in theorem 2, other bidders gain from not having to consider the possibility that additional bidders might represent bidding clubs. Paradoxically, other bidders' gain from the existence of a given bidding club is greater than the gain of that club's members.

**COROLLARY 2.** *In the equilibrium described in theorem 1, singleton bidders and members of other bidding clubs have higher expected utility when other agents participate in a given bidding club of size  $k \geq 2$ , as compared to a case where  $k$  additional agents participate directly in the main auction as singleton bidders.*

*Proof.* Consider a singleton bidder in the first case, where the club of  $k$  agents does exist. (It is sufficient to consider singleton bidders, since other bidding clubs bid in the same way as singleton bidders.) Following the equilibrium from theorem 1 this agent would submit the bid  $b^e(v_i, P^{n,1})$ . Theorem 2 shows that it is better to belong to a bidding club (and thus to bid  $b^e(v_i, P^{n,k})$ ) than to be a singleton bidder in an auction with the same number of agents (and thus to bid  $b^e(v_i, P^{n+k-1,1})$ ). Since the distribution  $P^{n,k}$  is just  $P^{n,1}$  with  $k-1$  singleton agents added,  $\forall k \geq 2, b^e(v_i, P^{n,1}) < b^e(v_i, P^{n,k})$ . Thus  $\forall k \geq 2, b^e(v_i, P^{n,1}) < b^e(v_i, P^{n+k-1,1})$ . ■

Finally, we can show that agents are indifferent between participating in the equilibrium from theorem 1 in a bidding club of size  $k$  (thus, where the number of agents is distributed according to  $P^{n,k}$ ) and participating in an economic environment with a stochastic number of bidders distributed according to  $P^{n,k}$ , but with no coordinators.

**THEOREM 3.** *For all  $\tau_i \in \mathcal{T}$ , for all  $k \geq 1$ , for all  $n \geq 2$ , agent  $i$  obtains the same expected utility by:*

1. *participating in a bidding club of size  $k$  in the economic environment from section 5.1 and following the equilibrium from theorem 1;*
2. *participating in a first-price auction with participation revelation in an economic environment with a stochastic number of bidders distributed according to  $P^{n,k}$  where all bidders receive the null signal, and where there are no coordinators.*

*Proof.* First we will show that agent  $i$ 's expected utility in case (2) above is the same as in a classical first-price auction with a stochastic number of bidders (i.e., without participation revelation). Second, we will show that agent  $i$ 's expected utility in this classical stochastic setting is the same as in case (1) above.

From proposition 4 it is an equilibrium for agent  $i$  to bid  $b^e(v_i, j)$  in a first-price auction with participation revelation (case (2)), where  $j$  is the number of bidders announced by the auctioneer. Since the number of agents is distributed according to  $P^{n,k}$ , the expected payment of agent  $i$  is  $\sum_{j=2}^{\infty} p_j^{n,k} b^e(v_i, j)$ . This is the definition of  $b^e(v_i, P^{n,k})$  from equation 4.

From proposition 3 this is an equilibrium bid of agent  $i$  when the number of agents is distributed according to  $P^{n,k}$  (without information revelation). Since both the classical first-price auction with a stochastic number of bidders and the first-price auction with participation revelation are efficient, agent  $i$ 's expected utility is the same under both auctions.

Under the equilibrium from theorem 1 (case (1)) the amount of  $i$ 's payment will be  $b^e(v_i, P^{n,k})$  if he wins. Since both the mechanism from case (1) and the classical first-price auction with a stochastic number of bidders are efficient, agent  $i$  has the same expected utility in both auctions. ■

This theorem shows that an agent would be as happy in a world without bidding clubs as he is in our economic environment. The difference between the two worlds is that in the latter bidding club coordinators make a positive profit on expectation, and indeed never lose money. That is, in the bidding club economic environment some expected profit is shifted from the auctioneer to the bidding club coordinator(s) without affecting the bidders' expected utility. We observe that it would be easy for coordinators to redistribute some of these gains to bidders along the lines of the second-price auction protocol proposed by Graham and Marshall: coordinators make a payment to every bidder who accepts the invitation to join, where the amount of this payment is less than or equal to the *ex ante* expected difference that bidder makes to the coordinator's profit. With this modification coordinators would be budget balanced only on expectation (violating requirement 2 from section 1.3), but agents would strictly prefer the bidding club economic environment to the economic environment in which coordinators are not present.

## 7. DISCUSSION

In this section we consider the trustworthiness and legality of coordinators, and also discuss two ways for auctioneers to disrupt bidding clubs in their auctions.

### 7.1. Trust

Why would a bidding club coordinator be willing to provide reliable service, and likewise why would bidders have reason to trust a coordinator? For example, a malicious coordination protocol could be used simply to drop all its members from the auction and reduce competition. While this is a reasonable concern, all the bidding club protocols discussed in this paper allow the coordinator to make a profit on expectation. There is thus incentive for a trusted third party to run a reliable coordination service. Indeed, coordinators would be very inexpensive to run: as their behavior is entirely specified, they could operate without any human supervision. The establishment of trust is exogenous to our model; we have simply assumed that all agents trust coordinators and that all coordinators are honest.

## 7.2. Legality

We have often been asked about the legal issues surrounding the use of bidding clubs. While this is an interesting and pertinent question, it exceeds both our expertise and the scope of this paper. We should note, however, that uses of bidding clubs exist that might not fall under the legal definition of collusion. For example, a corporation could use a bidding club to choose one of its departments to bid in an external auction. In this way the corporation could be sure to avoid bidding against itself in the external auction while avoiding dictatorship and respecting each department's self-interest. Coordinators may also be permitted by the auctioneer: e.g., by an internet market seeking to attract more bidders to its site.

## 7.3. Disrupting Bidding Clubs

There are two things an auctioneer can do to disrupt bidding clubs in a first-price auction. First, she can permit "false-name bidding." Our auction model has assumed that each agent may place only a single bid in the auction, and that the center has a way of uniquely identifying agents. For example, the auctioneer might use user accounts keyed to credit card billing addresses in combination with a reputation ranking, making it impossible for bidders to place bids claiming to originate from different agents. Second, she can refrain from publicly disclosing the winner of the auction.

If bidders can bid both in their bidding clubs and in the main auction, they are better off deviating from the equilibrium in theorem 1 in the following way. A bidder  $i$  can accept the invitation to join the bidding club but place a very low bid with the coordinator; at the same time,  $i$  can directly submit a competitive bid in the main auction. Agent  $i$  will gain by following this strategy when all other agents follow the strategies specified in theorem 1 because accepting the invitation to join the bidding club ensures that the club does drop all but one of its members and also causes the high bidder to bid less than he would if he were not bound to the coordination protocol. If the bidding club drops any bidders other than  $i$  then all agents' bids will also be lowered because the number of participants announced by the auctioneer will be smaller, compared to the case where the bidding club did not exist or where it was disbanded. However, if false-name bidding is impossible and the winner of the auction is publicly disclosed then the bidding club coordinator can detect an agent who has deviated in this way. Because the agent has agreed to participate in the bidding club the coordinator has the power to impose a punitive fine on this agent, making the deviation unprofitable. If either or both of these requirements does not hold, however, the coordinator will be unable to detect defection and so the equilibrium from theorem 1 will not hold.

## 8. CONCLUSION

We have presented a formal model of bidding clubs which departs in many ways from models traditionally used in the study of collusion; most importantly, all agents behave strategically based on correct information about the economic environment, including the possibility that other agents will collude. Other features of our setting include a stochastic number of agents and a stochastic number of bidding clubs in each auction. Agents' strategy space is expanded so that the decision of whether or not to join a bidding club is part of an agent's choice of strategy. Bidding clubs never lose money, and gain on expectation. We have showed a bidding club protocol for first-price auctions that leads to a (globally) efficient allocation in equilibrium, and which does not make use of side-payments. There are three ways of asking the question of whether agents gain by participating in bidding clubs in first-price auctions:

1. Could any agent gain by deviating from the protocol?
2. Would any agent be better off if his bidding club did not exist?
3. Would any agent would be better off in an economic environment that did not include bidding clubs at all?

We have showed that agents are strictly better off in the first two senses and no worse off in the last sense; furthermore, we have described a simple side-payment scheme that would make agents strictly better off in all three senses. We have also showed that each bidding club causes *non-members* to gain in the second sense. Finally, we have discussed ways for an auctioneer to set up the rules of her auction so as to disrupt bidding clubs.

## REFERENCES

- Cramton, P. and Palfrey, T. (1990). Cartel enforcement with uncertainty about costs. *International Economic Review*, 31(1):17–47.
- Feinstein, J., Block, M., and Nold, F. (1985). Assymetric behavior and collusive behavior in auction markets. *The American Economic Review*, 75(3):441–460.
- Graham, D. and Marshall, R. (1987). Collusive bidder behavior at single-object second-price and english auctions. *Journal of Political Economy*, 95:579–599.
- Graham, D., Marshall, R., and Richard, J.-F. (1990). Differential payments within a bidder coalition and the shapley value. *The American Economic Review*, 80(3):493–510.

- Harstad, R., Kagel, J., and Levin, D. (1990). Equilibrium bid functions for auctions with an uncertain number of bidders. *Economic Letters*, 33(1):35–40.
- Hendricks, K. and Porter, R. (1989). Collusion in auctions. *Annale's D'economie de Statistique*, 15/16:216–229.
- Leyton-Brown, K., Shoham, Y., and Tennenholtz, M. (2000). Bidding clubs: institutionalized collusion in auctions. In *ACM Conference on Electronic Commerce*.
- Mailath, G. and Zemsky, P. (1991). Collusion in second-price auctions with heterogeneous bidders. *Games and Economic Behavior*, 3:467–486.
- McAfee, R. and McMillan, J. (1987). Auctions with a stochastic number of bidders. *Journal of Economic Theory*, 43:1–19.
- McAfee, R. and McMillan, J. (1992). Bidding rings. *The American Economic Theory*, 82:579–599.
- Monderer, D. and Tennenholtz, M. (2000). Optimal Auctions Revisited. *Artificial Intelligence*, 120(1):29–42.
- Riley, J. and Samuelson, W. (1981). Optimal auctions. *American Economic Review*, 71:381–392.
- Robinson, M. (1985). Collusion and the choice of auction. *Rand Journal of Economics*, 16(1):141–145.
- von Ungern-Sternberg, T. (1988). Cartel stability in sealed bid second price auctions. *The Journal of Industrial Economics*, 18(3):351–358.

## Appendix G

# Mechanism Design with Incomplete Languages

Amir Ronen\*

Department of Computer Science  
Stanford University  
(amirr@Robotics.Stanford.EDU)

### Abstract

A major achievement of mechanism design theory is the family of truthful mechanisms often called VCG (named after Vickrey, Clarke and Groves). Although these mechanisms have many appealing properties, their essential intractability prevents them from being applied to complex problems like combinatorial auctions. In particular, VCG mechanisms require the agents to **fully describe** their valuation functions to the mechanism. Such a description may require exponential size and thus be infeasible for the agents.

A natural approach for this problem is to introduce an intermediate language for the description of the valuations. Such a language must be succinct to both the agents and the mechanism. Unfortunately, the resulting mechanisms are neither truthful nor do they satisfy individual rationality.

This paper suggests a general method for overcoming this difficulty. Given an intermediate language and an algorithm for computing the results, we propose three different mechanisms, each more powerful than its predecessor, but also more time consuming. Under reasonable assumptions, the results of our mechanisms are **at least as good** as the results of the algorithm on the actual valuations. All of our mechanisms have polynomial computational time and satisfy individual rationality.

## 1 Introduction

### 1.1 Motivation

The theory of mechanism design may be described as studying the design of protocols under the assumption that the participants behave according

---

\*This research was supported by Darpa grants number F30602-98-C-0214 and F30602-00-2-0598.

to their own goals and preferences and not necessarily as instructed by the protocol. The canonical mechanism design problem can be described as follows: A set of rational agents need to collaboratively choose an outcome  $o$  from a finite set  $O$  of possibilities. Each agent  $i$  has a privately known valuation function  $v^i : O \rightarrow R$  quantifying the agent's benefit from each possible outcome. The agents are supposed to report their valuation functions  $v^i(\cdot)$  to some centralized mechanism that chooses an outcome  $o$  that maximizes the total welfare  $\sum_i v^i(o)$ . The main difficulty is that agents may choose not to reveal their true valuations but rather report carefully designed lies in an attempt to influence the outcome to their liking. The tool that the mechanism uses to motivate the agents to reveal the truth is monetary payments. These payments are to be designed in a way that ensures that rational agents always reveal their true valuations – making the mechanism, so called, incentive compatible or truthful. To date there is only one general technique known for designing such a payment structure, sometimes called the generalized Vickrey auction [21], the Clarke pivot rule [1] the Groves mechanism [5], or, as we will, VCG. In certain senses this payment structure is unique [4, 17].

Although VCG mechanisms have many appealing properties, their **intractability** prevents them from being applied to complex problems like combinatorial auctions. This intractability is twofold: Firstly, VCG mechanisms require the agents to fully describe their valuation functions. Secondly, it requires the mechanism to find the optimal allocation.

The problem of combinatorial auctions (CA) is an important example of a mechanism design problem. In CA, the designer would like to auction a set  $S$  of items (e.g. radio spectra licenses) among a group of agents who desire them. As items may be substitutes (e.g. two licenses in the same place) or complementary (e.g. licenses in two neighboring states) the valuation of each agent may have a complex structure. A formal definition of the

problem can be found in section 2.1.

Consider a VCG mechanism for CA: The mechanism first asks each agent to declare her valuation function, i.e. to report a function  $w^i : 2^S \rightarrow R_+$ . It then computes the optimal allocation and the payments of each agent.

Such a mechanism is clearly intractable. Firstly, finding the optimal allocation is NP-hard even to approximate. Secondly, the mechanism relies on the agents' ability to describe their valuations in a way which is succinct to its allocation algorithm. This ability cannot be taken for granted. For example a naive solution will require each agent to report a vector of  $2^{|S|} - 1$  numbers to the mechanism. This of course is not feasible unless the number of items is very small. On the other extreme the designer can ask the agents to submit *oracles*, i.e. programs that return for every set  $s$  their valuation  $v^i(s)$ . However, it is not difficult to see that in order to find the optimal allocation or even a reasonable one, the allocation algorithm must query these oracles an exponential number of times. The natural solution for this problem is to introduce the notion of a *bidding language*. Such a language should enable the agents to efficiently represent or at least to approximate their valuations, but should also allow the allocation algorithm to compute the desired allocation in polynomial time. Hopefully such a language will capture most "real life" valuations. Various bidding languages were proposed in recent years. The interested reader is pointed to [11].

The drawback of this approach is that there are always valuation functions which are impossible to represent in polynomial-time. We therefore call such languages incomplete. Since VCG mechanisms with incomplete languages are not optimal, the impossibility results of [13] imply that they *cannot* be truthful! In other words, instead of describing their true valuation according to the designer's instructions, agents may have incentive to misreport. Therefore, there is no guarantee, even when the agents are rational, that the mechanism will find a reasonable allocation. Moreover, such

mechanisms do not even guarantee individual rationality. That is, there are cases where truthful agents will pay for their allocated sets more than their actual valuations for them.

Our goal in this paper is to prevent these phenomena.

## 1.2 This work

This paper proposes a general method for overcoming the non truthfulness of VCG mechanisms with incomplete languages. We first introduce the notions of oracles<sup>1</sup>, descriptions and consistency checkers in the context of VCG mechanisms. Oracles are programs that represent the agents' valuations. They are used by the mechanism to measure the agents' welfare. A consistency checker is a function that checks whether an agent's description, which is given in the intermediate language, is consistent with her oracle. These additions to the VCG method still do not suffice to guarantee its truthfulness.

We then describe three mechanisms which guarantee that under reasonable assumptions, truth-telling is *the* rational strategy for the agents. Each mechanism is more powerful but also more time consuming than its predecessor. All of our mechanisms have polynomial computational time.

Following [13] we adopt the concept of feasibly dominant actions (FDAs). Informally speaking, we assume that the agents choose their actions (strategies) according to their strategic knowledge. We say that an action is *feasibly dominant* if the agent is not aware of any circumstances where another strategy is better for her. It was argued in [13] that when feasibly dominant actions are available for the agent, it is irrational for her not to choose one of them. It was also shown in [13] that if the payment of a non-optimal mechanism is calculated according to the VCG formula, the existence of FDAs *must* rely on further assumptions on the agent's knowledge. Our mecha-

---

<sup>1</sup>Some advantages of using oracles were discussed in [19]

nisms guarantee that, under such reasonable assumptions, truth-telling is indeed an FDA. Each of them handles a more general form of knowledge than its predecessor (i.e. more sophisticated agents).

When the agents are truthful, the result of our mechanisms is **at least as good** as the result of the allocation algorithm on the truthfully reported descriptions. Our mechanisms also satisfy individual rationality.

Note that our method does not make any assumptions on the algorithm or the bidding language. The designer needs to design an intermediate language, a consistency checker and an allocation algorithm such that, when the agents prepare their descriptions according to her instructions, the overall result is good. She then gets the mechanism for free.

For simplicity we prove all our theorems directly for the combinatorial auction problem. Our results however are much more general and can be applied to any VCG, weighted VCG or compensation and bonus [14] mechanism.

### 1.3 Related work

Non optimal VCG mechanisms were first studied in [13]. This paper discusses VCG mechanisms where the optimal algorithm is replaced by a poly-time approximation or heuristic. This paper shows that mechanisms constructed this way *cannot* be truthful. It then proposes a general way of dealing with this non-truthfulness using a certain form of appeal functions.

The problem of combinatorial auctions has been studied by several researchers in recent years. A comprehensive survey of various aspects of this problem can be found in [2]. In particular, various bidding languages [3, 7, 20] and restrictions on the classes of bids that can be submitted (e.g. [6]) were proposed. A comparative study of some of these languages can be found in [11].

An alternative approach to the one that is taken here is to consider

mechanisms where the agents are not required to declare their valuation functions (non-revelation mechanisms). Examples of such mechanisms are the simultaneous ascending auction [10] and iBundle [16]. The efficiency of these auctions however is dependent on strong assumptions on the agents' behaviour. They are also specifically designed to address the combinatorial auction problem.

Finally, there is an extensive literature in the field of mechanism design. An introduction can be found in [8, chapter 23] and [15, chapter 10]

**Organization of this paper:** The rest of the paper is organized as follows: Section 2 formally defines combinatorial auctions and VCG mechanisms for CA and explains their intractability. Section 3 provides an example of a VCG mechanism with incomplete language and demonstrates the drawbacks of such mechanisms. Section 4 defines our most basic mechanism, describes the main concepts of [13] and shows that under reasonable assumptions on the agents' knowledge, truth-telling is an FDA. Sections 4 to 6 define extended versions of this mechanism and prove their basic properties. Section 7 discusses additional implementation issues and section 8 concludes the paper.

## 2 Preliminaries

### 2.1 Combinatorial auctions (CA)

The problem of combinatorial auctions (CA) has been extensively studied in recent years (see e.g. [7] [20] [3] [6] [11] ). The importance of this problem is twofold. Firstly, several important applications rely on it (e.g. the FCC auction [9]). Secondly, it is a generalization of many other problems of interest, in particular in the field of electronic commerce. A recent survey of various aspects of this problem can be found in [2]. For simplicity we prove all our theorems directly for this problem.

**The problem:** A seller wishes to sell a set  $S$  of items (radio spectra licenses, electronic devices, etc.) to a group of  $n$  agents who desire them. Each agent  $i$  has, for every subset  $s \subseteq S$  of the items, a non-negative number  $v^i(s)$  that represents how much  $s$  is worth for her. The function  $v^i(\cdot)$  is called the agent's *valuation* or *type*. We assume that  $v^i(\cdot)$  is **privately** known to the agent. Given a (possibly partial) allocation  $s = (s^1, \dots, s^n)$  we shall define the *total welfare* of the agents as  $g = \sum_i v^i(s)$ . In this paper we will be interested in mechanisms (protocols) which are designed to *maximize* the total welfare. This goal is justified in many settings. There is also a basic correlation between maximizing welfare and maximizing the seller's revenue. Solving the problem without monetary transfers is impossible (see a discussion at [8, chapter 23]). We assume that the mechanism can ask for payment from the agents and that the overall *utility* of each agent  $i$  is  $u^i = v^i(s) + p^i$  where  $s$  denotes the chosen allocation and  $p^i$  the amount of currency that the mechanism pays to the agent<sup>2</sup>. In an auction,  $p^i$  will be non-positive. This utility is what each *agent* tries to maximize.

For the sake of the example we take some standard additional assumptions on the type space of the agents:

**No externalities** The valuation of each agent depends only on the items allocated to her. I.e.  $\{v^i(s^i) | s \subseteq S\}$  completely represents the agent's valuation.

**Free disposal** Items have non-negative values. I.e if  $s \subseteq t$  then  $v^i(s) \leq v^i(t)$ .

**Normalization**  $v^i(\emptyset) = 0$ .

Note that the problem allows items to be complementary, i.e.  $v^i(S \cup T) \geq v^i(S) + v^i(T)$  or substitutes, i.e.  $v^i(S \cup T) \leq v^i(S) + v^i(T)$

---

<sup>2</sup>This is called the quasi-linearity assumption.

$(S, T$  disjointed). For example an agent may be willing to pay \$200 for a TV set, \$150 for a VCR, \$450 for both and only \$200 for two VCRs. The structure of the valuation functions might therefore be complex. The problem of finding an optimal allocation is equivalent to set-packing and is  $NP$ -hard even to approximate within any reasonable factor.

Note that the valuation functions are not known to the mechanism in advance. Moreover, if the mechanism is not carefully designed, the agents will have an incentive to **manipulate** it for their own self interest. Such manipulations might severely damage the efficiency of the mechanism. In mechanism design problems the agents are assumed to be *rational* in a game theoretic sense. They choose strategies which are good for them and not necessarily act as instructed. The goal of the designer is to design a mechanism (protocol) that produces good results under this assumption. Comprehensive surveys of mechanism design theory can be found in [15, chapter 10] [8, chapter 23].

In order to handle complex problems like combinatorial auctions the mechanism needs to address the following issues:

- Agents' valuations might be complex to express.
- The allocation and payments might be hard to compute.
- The mechanism needs to be designed to find good allocations even though the agents follow their own self interest.

Let us summarize our notations and terminology regarding this problem.

**Notations:** We shall denote the whole set of items by  $S$  and a (possibly partial) allocation by  $s = (s^1, \dots, s^n)$ . Note that the  $s^i$ 's are disjointed. We denote the type of agent  $i$  by  $v^i$  and the group's type by  $v = (v^1, \dots, v^n)$ . Let  $p^i$  denote the amount of currency that the mechanism pays to each agent  $i$  and  $u^i$  the agent's utility. Given an allocation  $s$  and a type  $v$  we denote by

$g_s(v)$  the welfare  $\sum_i v^i(s^i)$ . Finally we shall use the following vectorial notation: given a vector  $a = (a^1, \dots, a^n)$  we let  $a^{-i} = (a^1, \dots, a^{i-1}, a^{i+1}, \dots, a^n)$  and  $(b^i, a^{-i})$  denote the vector  $(a^1, \dots, a^{i-1}, b^i, a^{i+1}, \dots, a^n)$ .

## 2.2 VCG mechanisms for CA

One of the major achievements of mechanism design theory is the VCG method for constructing truthful mechanisms. In this subsection we briefly describe these mechanisms for CA and discuss some of their properties.

The simplest kind of mechanisms are protocols (called revelation mechanisms) where the agents are simply required to (privately) report their types to the mechanism. According to these declarations the mechanism computes the allocation and the payments. Note that agents may **lie** if it is beneficial for them. Such a mechanism can be denoted by a pair  $m = (k(w), p(w))$  where  $k$  denotes the allocation function,  $p$  the payment function and  $w$  the agents' declaration .

**Definition 1 (truthful mechanism)** *A revelation mechanism is called truthful if truth-telling is a dominant strategy for all agents. I.e. if lying to the mechanism can never be more beneficial than declaring  $v^i$ .*

VCG mechanism are a special kind of revelation mechanisms.

**Definition 2 (VCG mechanism)** *A VCG mechanism for CA is a revelation mechanism  $m = (k(w), p(w))$  such that:*

- *The mechanism chooses an allocation  $s = k(w)$  that maximizes the total welfare  $g_s(w)$  according to the declaration  $w$ .*
- *The payment is calculated according to the VCG formula:  $p^i(w) = \sum_{j \neq i} w^j(s) + h^i(w^{-i})$  ( $h^i(\cdot)$  can be any real function of  $w^{-i}$ ).*

**Theorem 2.1** ([5]) *A VCG mechanism is truthful.*

**Proof:** Assume by contradiction that the mechanism is not truthful. Then there exists an agent  $i$  of type  $v^i$ , a type declaration  $w^{-i}$  for the other agents, and  $w^i \neq v^i$  such that  $v^i(k((v^i, w^{-i}))) + p^i((v^i, w^{-i})) + h^i(w^{-i}) < v^i(k((w^i, w^{-i}))) + p^i((w^i, w^{-i})) + h^i(w^{-i})$ . Let  $s = k((v^i, w^{-i}))$  denote the chosen allocation when the agent is truthful and let  $s' = k((w^i, w^{-i}))$ . The above inequality implies that  $g_s((v^i, w^{-i})) < g_{s'}((v^i, w^{-i}))$ . This contradicts the *optimality* of  $k(\cdot)$ .  $\square$

Rational agents will therefore reveal their true type to the mechanism. Thus, when agents are rational the mechanism will result in the *optimal* allocation!

Note that the main trick of this method is to identify the utility of truthful agents with the declared total welfare. Similar techniques were introduced in [14] for handling different type of problems. The results presented here are applicable to their methods as well.

Another desirable property of mechanisms is called *individual rationality*. This means that the utility of a truthful agent is guaranteed to be non-negative. A special kind of VCG mechanism called Clarke's mechanism [1] can guarantee this property. It also guarantees that the payment of agents who are not allocated any object is zero. It does so by setting  $h^i = -\sum_{j \neq i} k(w^{-i})$  where  $k(w^{-i})$  denotes the result of the algorithm when agent  $i$  is "ignored". Until section 7 we shall only be interested in truthfulness. Thus, for simplicity we can assume that  $h^i(w^{-i}) \equiv 0$ .

It is worth notifying that weighted VCG mechanisms are possible as well (see e.g. [17] [14]). Also the designer can impose her own preferences by "pretending" to be one of the agents. To date VCG is the **only** general known method for the construction of truthful mechanisms. There is also some evidence [17] that other methods are generally impossible.

### 2.3 The intractability of VCG mechanisms

Although VCG mechanisms have many desirable properties, their essential intractability prevents them from being used for complex problems like CA. This intractability is twofold: VCG mechanisms require the agents to **fully describe** their valuation functions and require the mechanism to find optimal allocations.

The second aspect has been extensively discussed in [13]. This paper discusses VCG mechanisms where the optimal algorithm is replaced by a poly-time approximation or heuristics. It shows that mechanisms which are constructed in this way *cannot* be truthful. The paper proposes a method to overcome this non-truthfulness. It suggests a bounded rationality variant of truthfulness called feasible truthfulness and shows that under reasonable assumptions there is a general way of constructing poly-time feasible truthful mechanisms.

An even more fundamental obstacle on the way to the application of VCG mechanisms (and revelation mechanisms in general) to complex problems is the fact that the agents are required to describe their valuation functions to the mechanism. Consider for example a VCG mechanism for CA. One natural way in which an agent can describe her valuation function to the mechanism is by reporting a vector of numbers denoting her valuation for every possible combination of items. This however is infeasible unless the number of items is very small as it will require a vector of size  $2^{|S|} - 1$ . On the other extreme, the designer can ask the agent to construct an *oracle*, i.e. a program that returns for every set  $s$  the agent's valuation  $v^i(s)$ . However it is not difficult to see that in order to find the best allocation or even a reasonable one, the algorithm needs to query the oracle an exponential number of times.

The natural solution for this problem is to introduce the notion of a

*bidding language* (see e.g. [11]) – a language that will enable agents to efficiently represent or at least approximate their valuations but will also allow the mechanism’s algorithm to compute the desired allocation in polynomial time. Hopefully such a language will capture most ”real life” valuations. In addition the designer must provide the agents with *instructions* of how to construct these descriptions from their actual valuations. Given such a language  $L$  we can define VCG mechanisms as before. The bidding language and allocation algorithm must be constructed in a way that when the agents follow the designer’s instructions, the results will be good (heuristically, within a certain factor from the optimum etc.)

The problem with this approach is that there are always valuation functions which are *impossible* to represent in polynomial-time. We therefore call such languages incomplete. As such a mechanism is not optimal, the impossibility results in [13] imply that VCG mechanisms with incomplete languages *cannot* be truthful! In other words, agents may have incentives not to follow the designer’s instructions. Therefore there is no guarantee, even when the agents are rational, that the overall results will be good. Moreover, such mechanisms do not even guarantee individual rationality. That is, there are cases where truthful agents will pay for their allocated sets more than their actual valuations for them.

In this paper we propose a general method for overcoming this non-truthfulness. Our solution is in the same spirit of [13]. However several additional steps are needed to guarantee the good game theoretical properties of the resulting mechanisms.

### 3 Example VCG with OR bids

In this section we describe a simple example for a VCG mechanism with an incomplete bidding language. We shall use this example throughout

	A	B	AB
Agent1:	1	1	1.25 (2)
Agent2:	0.8	0.8	1.2 (1.6)

Figure 1: Type matrix for the OR example

the paper. We first describe the language and the mechanism. Then we analyze what strategies rational agents might choose when participating in it. Note that our language is less expressive than what we expect from real life mechanisms. We will demonstrate that even with such a language it is possible to construct mechanisms where truth-telling is *the* rational strategy.

Following [11] we define an *atomic bid* to be a pair  $(s, p)$  where  $s \subseteq S$  is a set of items and  $p$  is a price. The semantic of such a bid is "my maximum willingness to pay for  $s$  is  $p$ ". A description in this language consists of a *polynomial* number of such pairs. Given such a description  $(s_j, p_j)$  we can define, for every set  $s$ , the price  $p_s$  to be the maximal<sup>3</sup> sum of  $p_j$ s such that  $s_j \subseteq s$  are disjointed:  $\max\{\sum_j p_j | (s_j \subseteq s) \text{ and } \forall j \neq k, s_j \cap s_k = \emptyset\}$ . This so called OR language was used in [20].

**Proposition 3.1** [11] *OR bids can represent only super-additive valuation functions.*

□

The OR language therefore assumes that if an agent is willing to pay up to  $P_A$  for item  $A$  and  $P_B$  for item  $B$ , then she is willing to pay at least  $(P_A + P_B)$  for both.

Consider now the following (toy) example of a VCG mechanism: There are only two items  $A$  and  $B$ . As shown in figure 3, the type of Agent1 is  $(1, 1, 1.25)$  and of Agent2 is  $(0.8, 0.8, 1.2)$ .

---

<sup>3</sup>For the sake of the example we ignore the fact that computing this maximum might be *NP*-hard.

Suppose that the designer instructs the agents to submit their true valuation for every singleton. In this case we can define a description  $d^i = \{(s_j, p_j)\}$  as truthful if for every item  $j$ ,  $p_j = v^i(j)$ . In other words, such a description was prepared according to the designer's instructions. Consider a VCG mechanism with this language. After the descriptions are reported, the mechanism allocates the items optimally (according to the descriptions but not to the actual allocations!). It then calculates the payments according to the VCG formula. We assume that the designer has a small reserved price for each item, so objects which are not desired by the agents are not allocated.

In the example, when both agents are truthful, the mechanism will assign the valuation in brackets to the set AB (see figure 3). The mechanism in this case will allocate both items to Agent1 resulting in a utility of  $u^i = 1.25$  for each agent (recall that we assume the simplified form where  $h^i \equiv 0$ ). The optimal allocation will allocate to each agent one item, resulting in a welfare of 1.8.

The above mechanism is **not** truthful. For example if Agent1 "gives up" item B and declares  $(1, 0, 1.25)$  while Agent2's declaration remains the same, it will cause the algorithm to produce the optimal result and therefore will increase Agent1's utility to 1.8! The same is true for Agent2. On the other hand if both agents are "giving up" the same item, only one item will be allocated (to Agent1). This will result in a welfare of only 1.0. We shall call a declaration where the agent reports a 0 value on one of the items *singleton concession*. Another reasonable strategy for an agents is to find a description which will bring the mechanism's interpretation as close as possible to her actual valuation. Formally we define the  $l_\infty$ -approximation of  $v^i(\cdot)$  to be the description that minimizes  $\max_s |v^i(s) - d^i(s)|^4$ . Such a

---

<sup>4</sup>For the sake of the example we ignore the fact that calculating such a description might be NP-hard.

description for Agent1 is  $\{2/3, 2/3, 4/3\}$ . Note that the worthwhileness of such declarations is highly dependent on the declarations of the others. There are cases where such declarations will considerably improve the result of the algorithm and therefore will increase the agent's utility. On the other hand there are many cases where such designated "lies" will severely damage the total welfare and henceforth the agent's utility.

Note that the Clarke version of the above mechanism does **not** satisfy individual rationality. For example, if both agents are truthful, Agent1 gets both items, but pays 1.6, thereby loosing 0.35.

In this paper we will try to prevent these bad phenomena from happening.

## 4 Mechanism1

In this section we describe our first and most basic mechanism. We first describe the building blocks of the mechanism – oracles, descriptions and consistency checkers. Then we define the mechanism and formulate its basic properties. Finally we show that under reasonable assumptions truth-telling is *the* rational strategy for the agents.

We start with a formal definition adopted from [13] of computationally bounded algorithms<sup>5</sup>.

**Definition 3 (algorithm of degree  $d$ )** *Let  $n$  denote the number of agents. We say that a function  $F$  is of degree  $d$  if its running time is bounded by some polynomial of degree  $d$  of  $n$ .*

Our mechanism fixes a constant  $c = O(n^d)$  and terminates each function that runs more than  $c$  time units (see section 7 for more details).

---

<sup>5</sup>There are several alternative definitions. This one simplifies the formalization of the results.

## 4.1 Oracles and valid descriptions

All the mechanisms described in this paper ask the agents to prepare oracles that represent their valuation functions. These oracles are queried by the mechanisms in order to measure the total welfare. Formally:

**Definition 4 (oracle)** *An oracle is a function  $w : 2^S \rightarrow R_+$ . It is called truthful for agent  $i$  if  $w^i(s) = v^i(s)$  for every set  $s$ .*

We shall assume that agents are capable of preparing such oracles<sup>6</sup>. We also assume that all the oracles are of degree  $d$ .

As mentioned earlier, it is hard for allocation algorithms to work with oracles. We assume that the allocation algorithm accepts as input descriptions in some bidding language (e.g. the OR language) and ask the agents to prepare such descriptions. A consistency checker verifies that the agents' descriptions are consistent with their oracles.

**Definition 5 (valid description)** *A consistency checker is a function  $\psi(w, d)$  such that:*

- *$\psi(w, d)$  gets an oracle  $w$  and a description  $d$  in the bidding language and returns a "corrected" oracle  $w'$ .*
- *for every oracle  $w$  there exists at least one description  $d$  such that  $w = \psi(w, d)$ . Such "fixpoint" descriptions are called valid.*

Semantically, a valid description was prepared according to the designer's instructions. Since the mechanism can always use the "corrected" oracle  $w'$  we shall assume that agents' descriptions are valid. We also assume that a consistency checker of degree  $d$  is available to the designer and that given a

---

<sup>6</sup>The tools which must be provided by the designer in order to make this assumption realistic are not discussed in this paper.

declaration  $d$ , the designer can compute an oracle  $w_d$  such that  $d$  is a valid description of  $w_d$ . We say that an agent's description is truthful if it is a valid description of a truthful oracle.

In the OR language for example we can define  $w'(s) = p$  for every atomic bid  $(s, p)$  in the description. Creating an oracle  $w$  from a description  $d$  such that  $d$  is valid is straight forward.

## 4.2 Appeal functions

Another basic building block of our mechanism is the notion of appeal functions. This is a modification of the appeals that were introduced in [13]. Intuitively an appeal function lets an agent incorporate her own *knowledge* about the algorithm into the mechanism. The idea is that instead of declaring a falsified type, the agent can follow the designer's instructions and ask the mechanism to *check* whether the false description would have lead to better results. The mechanism will then choose the better of these two possibilities leveraging both the agent's utility and the total welfare.

**Definition 6 (appeal)** *An appeal function gets as input the agents' oracles and valid descriptions and returns a tuple of alternative descriptions. I.e. it is of the form:  $l(w^1, \dots, w^n, d^1, \dots, d^n) = (d'^1, \dots, d'^n)$  where  $d^i$  is a valid description of  $w^i$ .*

Note that the  $d'^i$ 's do not have to be valid. The semantics of an appeal  $l$  is: "when the agents' type is  $w = (w^1, \dots, w^n)$  and is described by  $(d^1, \dots, d^n)$ , I believe that the output algorithm  $k$  produces a better result if it is given  $d'$  instead of the actual description  $d$ ".

We assume that all appeal functions are of degree  $d$  for some reasonable value of  $d$ . In section 7 we will discuss ways to enforce such a limit.

In our OR example (section 3) an appeal for Agent1 might try to give up one of the items (i.e. perform a singleton concession) or try to give up

item A for herself and B for Agent2 etc.

The actual implementation of the appeal functions is discussed in subsection 7.

### 4.3 Mechanism1

We can now define our first mechanism.

**Definition 7 (mechanism1)** *Given an allocation algorithm  $k(d)$ , and a consistency checker for the bidding language we define mechanism1 as follows:*

1. *Each agent submits to the mechanism:*
  - An oracle  $w^i(\cdot)$ .
  - A (valid) description  $d^i$ .
  - An appeal function  $l^i(\cdot)$ .
2. *Let  $w = (w^1, \dots, w^n)$ ,  $d = (d^1, \dots, d^n)$ . The mechanism computes the allocations  $k(d), k(l^1(w, d)), \dots, k(l^n(w, d))$  and chooses among these allocations the one that maximizes the total welfare (according to  $w!$ ). In other words, the mechanism tries all the appeals and chooses the one that yields the best result.*
3. *Let  $\hat{s}$  denote the chosen allocation. The mechanism calculates the payments according to the VCG formula:  $p^i = \sum_{j \neq i} w^j(\hat{s}) + h^i(w^{-i}, d^{-i}, l^{-i})$  ( $h^i(\cdot)$  can be any real function).*

Note that  $h^i(\cdot)$  is independent of agent  $i$ . Until section 7 we simply assume that it is always zero. Note also that we do not require the allocation algorithm  $k(\cdot)$  to be optimal. It can be *any* polynomial time approximation or heuristic.

An action (strategy) in mechanism1 is a triplet  $(w^i, d^i, l^i)$ . We say that such an action is *truthful* if  $w^i$  is truthful. The following two observations are key properties of the mechanism:

**Proposition 4.1** *Consider mechanism1 with an allocation algorithm  $k(\cdot)$ . Let  $d = (d^1, \dots, d^n)$  denote the agents' descriptions. If all the agents are truth-telling, the allocation chosen by the mechanism is at least as good as  $k(d)$ .*

□

**Proposition 4.2** *If the allocation algorithm  $k$ , the appeal functions, oracles and consistency checkers are of degree  $d$ , then the mechanism is of degree  $d + 2$ .*

□

Let  $\hat{s}$  denote the chosen allocation. Let  $\tilde{v} = (v^i, w^{-i})$ . Since we assume that  $h^i(\cdot) = 0$ , the utility of agent  $i$  equals  $g_{\hat{s}}(\tilde{v})$  – the total welfare when the allocation is  $\hat{s}$  and the type is  $\tilde{v}$ . Lying to the mechanism, i.e. submitting an oracle  $w^i \neq v^i$ , is thus beneficial for the agent *only* if it causes the mechanism to compute a better result (relatively to  $\tilde{v}$ ). (For a more comprehensive discussion see [13].) Note that when an agent lies to the mechanism, she may not only cause damage to the algorithm's result, but may also cause the mechanism to prefer the wrong allocation on the second stage. Thus, an agent needs to have a good *reason* for lying to the mechanism.

We will show that under reasonable assumptions on the agents, truth-telling is *the* rational strategy for the agents. Thus, when the agents are rational, the result of the mechanism is at least as good as the result of the allocation algorithm on the truthful descriptions.

#### 4.4 An example

Consider the OR example of section 3. Suppose that Agent1 notices that usually the result of the algorithm improves when she is giving up item A. In a VCG mechanism the agent may be tempted to misreport in order to increase the total welfare and henceforth her own utility. In many cases however this will cause damage to the overall welfare and henceforth to Agent1. In our mechanism Agent1 can, instead of lying, declare her **true** type to the mechanism and ask it to **check** whether such a lie would have been helpful. If so, it prefers the result that was obtained by "lying". Otherwise, the mechanism prefers the result of the algorithm on the truthful description and thus prevents the damage that would have been caused by the lie. This form of appeal functions provides the agents with a lot of power. Suppose, for example, that Agent1 notices that the result improves if she gives up item A while Agent2 is giving up item B. As before, the agent can ask the mechanism to check whether such a transformation of the input would have improve the overall result.

We note that not every knowledge of the agent about the allocation algorithm  $k(\cdot)$  can be exploited in this mechanism. Suppose that Agent1 notices that when both agents submit  $l_\infty$ -approximations of their valuations the overall result improves. However, as she is given an oracle for  $v^2$ , she cannot compute Agent2's approximation as it requires her to query the oracle for every possible subset. Therefore, she cannot exploit her knowledge about the algorithm. Such phenomena is problematic and do not occur in the setting of [13].

#### 4.5 When is it rational to tell the truth to the mechanism?

It was shown in [13] that even with full descriptions available, non-optimal VCG mechanisms cannot be truthful (unless they produce unreasonable re-

sults). That paper introduces a bounded rationality variant of the concept of dominant strategies called feasible dominance and shows that under reasonable assumptions truth-telling is feasibly dominant for the agents. This paper follows this pattern. In this section we first describe the basic concepts of [13]. We then consider mechanism1 and analyze the conditions under which truth-telling is feasibly dominant for the agents.

#### 4.5.1 Feasibly dominant actions (FDAs)

In this section we briefly describe the main concepts of [13]. The reader is referred to this paper for a more comprehensive discussion.

**Notations:** We denote the action (strategy) space of agent  $i$  by  $A^i$ . Given a tuple  $a = (a^1, \dots, a^n)$  of actions chosen by the agents, we denote the utility of agent  $i$  by  $u^i(a)$ .

In mechanism1 an action for the agent is a triplet  $(w^i, d^i, l^i)$ .

In classical game theory, given the actions of the other agents  $a^{-i}$ , the agent is (implicitly) assumed to be capable of responding by the optimal  $a^i$ . As the action space is typically very complex, this assumption is not natural in many real-life situations. The concept of feasibly dominant actions reformulates the concept of dominant actions under the assumption that the agent has only a limited capability of computing her response. It is meant to be used in the context of revelation games.

**Definition 8 (strategic knowledge)** Strategic knowledge (*or* response function) of agent  $i$  is a partial function  $b^i : A^{-i} \rightarrow A^i$ .

Knowledge is a function by which the agent describes (for herself!) how she would like to respond to any given situation. The semantics of  $a^i = b^i(a^{-i})$  is “when the others’ actions are  $a^{-i}$ , the best action which I can think of is  $a^i$ ”. The fact that  $a^{-i}$  is not in the domain of  $b^i$  means that the

agent does not know how to respond to  $a^{-i}$  or alternatively will not regret her choice of action when the others played  $a^{-i}$ . Naturally we assume that each agent is capable of computing her own knowledge and henceforth that  $b^i$  is of degree  $d$ .

**Definition 9 (feasible best response)** *An action  $a^i$  for agent  $i$  is called feasible best response to  $a^{-i}$  if either  $a^{-i}$  is not in the domain of the agent's knowledge  $b^i$  or  $u^i((b^i(a^{-i}), a^{-i})) \leq u^i(a)$ .*

In other words, other actions may be better against  $a^{-i}$  but at least when choosing her action the agent was not aware of these.

The definition of feasibly dominant actions now follows naturally.

**Definition 10 (feasibly dominant action)** *An action  $a^i$  for agent  $i$  is called feasibly dominant if it is a feasible best response against any  $a^{-i}$ . We also call such an action FDA .*

It was argued in [13] that if an agent has feasibly dominant actions available, then it is *irrational* not to choose one of them.

#### 4.5.2 When is it rational to tell the truth to the mechanism?

Recall that the overall utility of each agent  $i$  equals  $g_{\hat{s}}(\tilde{v})$  where  $\hat{s}$  denotes the chosen allocation and  $\tilde{v} = (v^i, w^{-i})$ . It is not difficult to see that when the agent declares a falsified valuation, there are cases where she will consequently lose. The agent needs therefore a good reason for lying to the mechanism. When the appeals of the agents are time-limited (i.e. of degree  $d$ ) it was shown in [13] that the existence of FDAs for the agents *must* rely on further assumptions on the agents' *knowledge*. Here we formulate two such assumptions and show how to construct computationally efficient truthful FDAs for the agents.

**Definition 11** ([13]) (**declaration based knowledge**) *Knowledge  $b^i(\cdot)$  is called declaration based if it is of the form  $b^i(w^{-i}, d^{-i}) = (w^i, d^i)$ .*

The semantics of declaration based knowledge is: "If I knew that the others declare  $(w^{-i}, d^{-i})$ , regardless of their appeals, I would like to declare  $(w^i, d^i)$ ". In our OR bids example of section 3, such knowledge for Agent1 may be: "If Agent2 has a high valuation for item B, I would like to give it up".

A declaration based knowledge naturally defines an appeal function which we also denote by  $b^i(\cdot)$ :  $b^i(w, d) = (b^i(w^{-i}, d^{-i}), d^{-i})$ .

**Theorem 4.3** *If  $b^i(\cdot)$  is a declaration based knowledge for agent  $i$  then  $(v^i, d^i, b^i)$  is feasibly dominant for the agent.*

**Proof:** Let  $\hat{s}$  denote the chosen allocation. Let  $\tilde{v} = (v^i, w^{-i})$ . Recall that the utility of agent  $i$  equals  $g_{\hat{s}}(\tilde{v})$ . Also let  $\phi$  denote the empty appeal. Assume by contradiction that there exists  $a^{-i} = (w^{-i}, d^{-i}, \phi^{-i})$  that contradicts the agent's knowledge. Note that the appeals of the other agents can be assumed empty and also that it must be that  $a^{-i}$  is in the domain of  $b^i(\cdot)$ . Let  $(w'^i, d'^i) = b^i(a^{-i})$ . Let  $s = k(d)$  and let  $s' = k(d'^i, d^{-i})$  denote the allocation when she lies. By the assumption,  $g_s(\tilde{v}) < g'_s(\tilde{v})$ . However when the agent truthfully submits  $(v^i, d^i, b^i)$  the mechanism computes  $s = k(d)$  and  $s' = k(d'^i, d^{-i})$  and takes the better among them according to  $\tilde{v}$ . A contradiction.  $\square$

**Definition 12** ([13])(**appeal independent knowledge**) *Knowledge  $b^i(\cdot)$  is called appeal independent if it is of the form  $b^i(w^{-i}, d^{-i}) = (w'^i, d'^i, l^i)$ .*

**Theorem 4.4** *If  $b^i(\cdot)$  is an appeal independent knowledge of agent  $i$  then there exists a truthful FDA of degree  $d$  for the agent.*

**Proof:** Define an appeal  $l^i$  as follows. Given  $(w^{-i}, d^{-i})$  let  $(w'^i, d'^i, l'^i) = b^i(w^{-i}, d^{-i})$ .  $l^i$  computes  $k(d'^i, d^{-i})$ ,  $k(l'^i((w'^i, w^{-i}), (d'^i, d^{-i})))$  and takes the best according to  $(v^i, w^{-i})$ . Since all the functions involved are of degree  $d$ , so is  $l^i(\cdot)$ . Similarly to theorem 4.3,  $(v^i, l^i)$  is an FDA.  $\square$

The semantics of declaration based knowledge is the same as declaration based except that the agent also submits an appeal  $l^i$ .

Agents who are not capable of reasoning about others' appeals or do not want to count on them would have appeal independent knowledge. We argue that this would be the most common case. In all of the examples of section 4.4 the agents' knowledge was appeal independent.

## 5 Mechanism2: moving information around

A major difficulty that arises when coping with incomplete languages is the asymmetric knowledge of the agents regarding their *own* valuations. For example, in the setting of section 3, it is reasonable to assume that Agent1 can compute her own  $l_\infty$ -approximation but Agent2 cannot compute it. Thus, Agent1 might face the following considerations:

- The result of the algorithm improves significantly when all agents report their  $l_\infty$ -approximations.
- Reporting my  $l_\infty$ -approximation instead of my truthful description, will enable Agent2 to compute the optimal result.

In other words, in mechanism1, agents may want to misreport in order to pass useful information about their own valuation to the others. In order to prevent this we modify the mechanism to allow the agents to convey such information.

**Definition 13 (information structure)** An information structure  $I^i$  for agent  $i$  is a sequence of descriptions (possibly with repetitions)

$(d_0, d_1, \dots, d_k)$  such that  $d_0$  is a valid description.

In addition we require each agent to provide for each  $d_j$  an *example*  $(w^{-i}, d^{-i})$  such that  $k(d_j, d^{-i})$  is a better allocation than  $k(d_0, d^{-i})$ . This is done in order to force the agents to submit only useful information.

$I^i$  contains additional information that the agent can pass to the others' appeals. The semantics of  $I_i$  is "My valid description is  $d_0$ . Nevertheless, I suggest that you first try to work with  $d_1$ , after that with  $d_2$ , etc". Many alternative ways to define such information structures are possible. It may be interesting to compare between different structures.

We can now define our second mechanism.

**Definition 14 (mechanism2)** *Given an allocation algorithm  $k(d)$ , and consistency checker for the bidding language we define mechanism2 as follows:*

1. *Each agent submits to the mechanism:*
  - *an oracle  $w^i$ . (let  $w = (w^1, \dots, w^n)$ )*
  - *an information structure  $I^i$ . (let  $I = (I^1, \dots, I^n)$ )*
  - *an appeal function of the form  $l^i(w, I) = d$ .*
2. *The mechanism computes the allocations  $k(d), k(l^1(w, I)), \dots, k(l^n(w, I))$  and chooses among these allocations the one that maximizes the declared total welfare.*
3. *The mechanism calculates the payments according to the VCG formula.*

We can now expand the definition of knowledge under which the existence of truthful FDAs is guaranteed. This definition refers to knowledge that was obtained by checking a representative family of (tuples of) appeals of the other agents.

**Definition 15** [13] ( **$d$ -obtainable knowledge**) *Knowledge  $b^i(\cdot)$  is called  $d$ -obtainable if the following holds:*

1.  $b^i$  is of degree  $d$ .
2. Every appeal function that appears in the domain or in the range of  $b^i(\cdot)$ , is of degree  $d$ .
3. There are at most  $n^d$  appeal functions that appear in the domain or in the range of  $b^i(\cdot)$ . Moreover there exists a representative family  $L^i$  of no more than  $n^d$   $(n - 1)$ -tuples of appeals such that for every tuple  $\varphi^{-i}$  that appears in the domain of  $b^i$  there exists a  $\psi^{-i} \in L^i$  such that for all  $(w^{-i}, I^{-i})$ ,  $b^i(((w^{-i}, I^{-i}), \varphi^{-i})) = b^i(((w^{-i}, I^{-i}), \psi^{-i}))$ .

The assumption that agents' knowledge is  $d$ -bounded is justified by the immense complexity of the appeal space. It assumes that an agent cannot think about more than a small family of representative cases  $L^i$ . For a more comprehensive discussion on this assumption see [13]. We need an additional assumption on the appeal class that the agent considers. We will remove this assumption later on.

**Definition 16 (monotonic appeal)** *We say that an appeal function  $l(\cdot)$  is monotonic if for every  $w$  and for every two structures  $I = (I^1, \dots, I^n)$  and  $I' = (I'^1, \dots, I'^n)$  such that  $I^j$  is a subset of  $I'_j$  for all  $j$ ,  $k(l(w, I'))$  is at least as good as  $k(l(w, I))$ .*

In other words, giving more information to the appeal can just help it to compute a better result. We cannot expect the appeals to be monotonic as such monotonicity usually requires exponential time. However, it is reasonable to think that appeals will be monotonic *in general*, that is that the addition of useful information and, in particular, of truthful descriptions,

usually helps the appeals to improve the overall result. Changing the order of the  $d_i$ s in the information structures does not affect monotonic appeals.

**Definition 17 (monotonic  $d$ -obtainable knowledge)** *Knowledge for agent  $i$  is called monotonic  $d$ -obtainable if it is  $d$ -obtainable and all the appeals that appear in its domain or in its range are monotonic.*

**Theorem 5.1** *If the agent's knowledge is monotonic  $d$ -obtainable, she has a truthful FDA of degree  $3 \cdot d$ .*

**Proof:** Let  $I^i$  denote a maximal sequence of useful information that an agent  $i$  can compute (i.e. it contains all the cases that the agent finds useful). Let  $b^i$  be a  $d$ -obtainable knowledge for agent  $i$ . Given  $(w^{-i}, I^{-i})$  we shall define an appeal  $l^i$  as follows: Let  $L$  be the family of all appeals that appear in the domain or in the range of  $b^i$ . Let  $L^i$  be the representative family. We define  $\omega$  to denote the set of all the "useful lies"  $\omega = \{w^i | \exists \psi^{-i} \in L^i, \varphi^i \text{ s.t. } (w^i, I^i, \varphi^i) = b^i(w^{-i}, I^{-i}, \psi^{-i})\}$ . Obviously  $|W|, |L|$  are bounded by a polynomial of degree  $d$ .

For every pair  $(w^i \in \omega, l \in L)$  we let  $l^i$  compute the result of  $l$  as if she had submitted  $(w^i, I^i, l)$ , i.e. compute  $k(l(w^i, w^{-i}), (I^i, I^{-i}))$ . The appeal returns the best of these allocations according to  $(v^i, w^{-i})$ .

As all the functions involved are of degree  $d$ , it is not difficult to verify that the appeal is of degree  $3 \cdot d$ .

We now show that submitting  $(v^i, I^i, L^i)$  is an FDA. Otherwise there exists a triplet  $(w^{-i}, I^{-i}, l^{-i})$  that contradicts  $b^i(\cdot)$ . Since  $b^i(\cdot)$  is  $d$ -obtainable we can assume that  $l^{-i}$  is in the representative family. Let  $(w^i, \iota^i, \delta^i) = b^i(w^{-i}, I^{-i}, l^{-i})$ . Because of the monotonicity we can assume that  $\iota^i$  contains all the useful information that  $i$  can think of (i.e.  $\iota^i = I^i$ ). However the appeal  $l^i$  checks the case where  $i$  submits  $(w^i, I^i, \delta^i)$ . Therefore  $l^i$ 's result must be at least as good as the result of the mechanism in this case – a contradiction.  $\square$

## 6 Mechanism3: adding meta-appeals

In section 5 we assumed that the agents' appeals are monotonic. Our final step is to get rid of this assumption. We first define the notion of a meta-appeal.

**Definition 18 (meta appeal)** *A meta appeal is a function that gets a vector of information structures  $I = (I^1, \dots, I^n)$  and returns a list of vectors of the form  $I' = (I'^1, \dots, I'^n)$  such that  $I'^j$  is a subset of  $I^j$ .*

In other words, the meta appeals compute a list of alternative information structures for the group. Note that many variants of this definition are possible. We assume that all the meta-appeals are of degree  $d$ .

**Definition 19 (mechanism3)** *Given an allocation algorithm  $k(d)$ , and a consistency checker for the bidding language we define mechanism3 as follows:*

1. *Each agent submits to the mechanism:*

- *An oracle  $w^i$ . (let  $w = (w^1, \dots, w^n)$ )*
- *An information structure  $I^i$ . (let  $I = (I^1, \dots, I^n)$ )*
- *An appeal function of the form  $l^i(w, I)$ .*
- *A meta appeal  $\chi^i(\cdot)$ .*

2. *The mechanism computes a list  $\Gamma$  containing all the results of the meta-appeals as well as the original tuple of information structures  $I$ .*

3. *The mechanism computes, for every pair  $(l^j, I')$  such that  $I' \in \Gamma$  and  $l^j$  is an appeal, the allocation  $k(l^j(w, I'))$ . It also computes  $k(d)$ . It then chooses among these allocations the one that maximizes the declared total welfare.*

4. The mechanism calculates the payments according to the VCG formula.

Note that the mechanism is of degree  $d + 3$ .

We can now define  $d$ -obtainable knowledge similarly to the previous section. We add however the condition that it ignores the meta appeals of the other agents.

**Definition 20** [13] ( **$d$ -obtainable knowledge of mechanism3**) *We say that knowledge  $b^i(\cdot)$  of mechanism3 is  $d$ -obtainable if the following holds:*

1.  $b^i(\cdot)$  is of degree  $d$ .
2.  $b^i(\cdot)$  ignores the meta-appeals of the other agents, i.e it is of the form  $b^i(w^{-i}, I^{-i}, l^{-i}) = (w^i, I^i, l^i)$ .
3. Every appeal function that appears, in the domain or in the range of  $b^i(\cdot)$ , is of degree  $d$ .
4. There are at most  $n^d$  appeal functions that appear in the domain or in the range of  $b^i(\cdot)$ . Moreover there exists a representative family  $L^i$  of no more than  $n^d$   $(n - 1)$ -tuples of appeals such that for every tuple  $\varphi^{-i}$  that appears in the domain of  $b^i(\cdot)$  there exists a  $\psi^{-i} \in L^i$  such that for all  $(w^{-i}, I^{-i})$ ,  $b^i(((w^{-i}, I^{-i}), \varphi^{-i})) = b^i(((w^{-i}, I^{-i}), \psi^{-i}))$ .

The main justification behind the assumption that  $b^i$  ignores the meta-appeals is that the space of meta-appeals is extremely complex. Moreover, properties of the meta-appeals are only partially connected to the actual bidding language or the algorithm. The only potential profit from lying that we can imagine are "extra-trials" of the allocation algorithm when the others' appeals are forced to use the agent's false description. We presume that such potential gains are negligible compared to the obvious loss caused by lying. It is also natural to think that if the appeals of the other agents

ignore the agent's recommendation to use  $d_1$ , they have a good reason to do so. We argue that knowledge which is not  $d$ -obtainable is unlikely to exist. However, this "thesis" needs to be checked *experimentally*.

**Theorem 6.1** *If the agent's knowledge is  $d$ -obtainable, then she has a truthful FDA  $(v^i, I, l^i, \chi^i)$  such that  $l^i$  is of degree  $3 \cdot d$  and  $\chi^i$  of degree  $d$ .*

**Proof:** Similarly to the proof of 5.1, given  $(w^{-i}, I^{-i})$ , we define the set of "useful lies"  $\omega = \{w^i | \exists \psi^{-i} \in L^i, \varphi^i \text{ s.t. } (w^i, I^i, \varphi^i) = b^i(w^{-i}, I^{-i}, \psi^{-i})\}$ , and the family  $L$  of appeals which appear in  $b^i(\cdot)$ . In addition we define the set of useful information structures  $\chi^i = \{I^i | \exists \psi^{-i} \in L^i, \varphi^i \text{ s.t. } (w^i, I^i, \varphi^i) = b^i(w^{-i}, I^{-i}, \psi^{-i})\}$ . We define  $I$  to be a union of all  $I \in \chi^i$ , an appeal  $l^i$  like in the proof of 5.1. The proof that  $(v^i, I, l^i, \chi^i)$  is an FDA is similar to 5.1.

□

### 6.1 Example: Mechanism3 with OR bids

Consider mechanism3 for CA with OR bids (section 3). Suppose that Agent1 notices the following phenomena:

1. When all agents perform  $l_\infty$ -approximations the result of  $k(\cdot)$  usually improves considerably.
2. The result also typically improves if agents perform singleton concessions on different items. The improvement however is less significant than in the first case.

Such an agent may anticipate three kinds of appeals:

- Appeals of agents that notice the first phenomenon and will therefore leverage from her  $l_\infty$ -approximation.
- Appeals of agents who notice only the second phenomenon and will only be disturbed by her  $l_\infty$ -approximation.

- Appeals that will work best with her valid description.

Mechanism3 gives Agent1 the possibility of constructing a strategy that will **dominate** every case that she can think of! She just needs to include in her meta appeal three information structures. One that includes only her valid description, one that will include her singleton concessions as well and one that will also include her  $l_\infty$ -approximation.

We note that there exist additional ways to justify why truth-telling is *the* rational strategy for the agents. Those are omitted from the paper mainly due to space constraints.

## 7 Other implementation issues

In this section we address two additional issues which a designer may face when implementing our mechanisms: guaranteeing individual rationality and forcing reasonable time limitations on the agents.

In [13] it was shown that the allocation algorithm can be transformed in polynomial time to an algorithm which satisfies additional monotonicity requirements. With such an algorithm it is possible to define the function  $h^i(\cdot)$  of our mechanisms similarly to Clarke's mechanism [1]. The proof that the resulting mechanisms satisfy individual rationality is similar to [13].

This paper shows that if enough computational time is given to the agents, they can construct truthful FDAs. On the other hand the mechanism needs to find a way to enforce *reasonable* time limits on the computational time of the agents, i.e. to enforce time limits on the appeals and meta appeals. This issue was discussed in [13]. In particular it was suggested that knowledge-reflecting structure will be chosen for description of the appeal functions. Such a structure enables the limitation of the computational time of the appeals according to the agents' own limitations and thus preserves the existence of truthful FDAs. We presume that severe limitations can

be imposed on the length of the lists produced by the meta appeals while preserving the existence of FDAs. Finally, we think that it is a good heuristic to charge small fees for extra computational time.

At a first glance our protocols may seem to put a lot of burden on the agents. However we argue that with the right tools (e.g. tools for building oracles), mechanisms like ours can become even more "agent friendly" than non-revelation mechanisms.

## 8 Conclusions and further research

In this paper we propose a general way to overcome the deficiencies of VCG mechanisms with incomplete languages. Given an intermediate language, a consistency checker, and an algorithm for the computation of the outcomes (e.g. allocations) we construct three mechanisms, each more powerful but also more time-consuming than its predecessor. All our mechanisms have polynomial computational time and satisfy individual rationality.

We adopt the strong concept of feasible dominant strategies of [13] which is a bounded rationality version of dominant strategies and showed that under reasonable assumptions on the agents' knowledge, truth-telling is feasibly dominant for the agents. In addition when an agent lies to the mechanism, there are cases where she will consequently lose.

When the agents are truth-telling the results of our mechanisms are **at least as good** as the mechanisms' algorithm. Our methods are general and can be applied to any VCG , weighted VCG or compensation and bonus [14] mechanism.

The paper assumes that in practice, agents will have only limited knowledge and thus will not be able to do better than their truthful FDAs. This thesis can and should be checked by experiments with "real" agents. On the other hand we feel that this assumption will remain true even when

severe time limitations are forced on the agents. In fact it will not even be a surprise if even in a VCG mechanism, if the bidding language and the allocation algorithm are reasonably designed, the agents will not be able to do better than truth-telling! This too can be checked experimentally.

Very little is currently known about the revenue of mechanisms for complex problems. In particular note that when a non-optimal VCG mechanism is naively used for a combinatorial auction, there are even cases where the mechanism must pay to the agents instead of vice-versa!

In our constructions, there are several tools that the designer must provide to the agents. Tools to construct oracles, descriptions, appeals etc. Methods for providing such tools were not discussed in this paper and are crucial for the success of our mechanisms.

Finally we note that it might be fruitful to explore the possibility of using appeal functions in situations where the agents have budget limits. When such limits exist, agents may have incentives to cause others to run out of budget and it is not likely that dominant strategy mechanisms exist. One natural way to deal with budget limits, is to truncate the agent's valuation to her limit and then use VCG [12]. Truth-telling in this mechanism is a *safe* strategy for the agent as she never pays more than her budget. We argue that appeals of certain forms can play the role of threats and prevent the worth-willingness of causing others to run out of budget.

**Acknowledgments:** We thank Yoav Shoham for many fruitful discussions. We thank Inbal Ronen for reviewing previous versions of this paper.

## References

- [1] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, pages 17–33, 1971.

- [2] Sven de Vries and Vohra Rakesh. Combinatorial auctions: A survey.  
To appear.
- [3] Yuzo Fujishima, Kevin Leyton-Brown, and Yoav Shoham. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *IJCAI-99*, 1999.
- [4] J. Green and J.J. Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, pages 427–438, 1977.
- [5] T. Groves. Incentives in teams. *Econometrica*, pages 617–631, 1973.
- [6] R. M. Harstad, Rothkopf M. H., and Pekec A. Computationally manageable combinatorial auctions. Technical Report 95-09, DIMACS, Rutgers university, 1995.
- [7] Daniel Lehmann, Liadan O‘Callaghan, and Yoav Shoham. Truth revelation in rapid, approximately efficient combinatorial auctions. In *ACM Conference on Electronic Commerce (EC-99)*, pages 96–102, November 1999.
- [8] A. Mas-Collel, Whinston W., and J. Green. *Microeconomic Theory*. Oxford university press, 1995.
- [9] J. McMillan. Selling spectrum rights. *Journal of Economic Perspectives*, pages 145–162, 1994.
- [10] Paul Milgrom. Putting auction theory to work: the simultaneous ascending auction. Technical Report 98-002, Dept. of Economics, Stanford University, 1998.

- [11] Noam Nisan. Bidding and allocation in combinatorial auctions. In *Proc. of the Second ACM Conference on Electronic Commerce (EC00)*, pages 1–12, 2000.
- [12] Noam Nisan. Private communication, 2000.
- [13] Noam Nisan and Amir Ronen. Computationally feasible vcg mechanisms. In *Proc. of the Second ACM Conference on Electronic Commerce (EC00)*, pages 242–252, October 2000.
- [14] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behaviour*, 2001. To appear; Extended abstract appeared in STOC99.
- [15] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT press, 1994.
- [16] David Parkes. ibundle: An efficient ascending price bundle auction. In *In Proc. ACM Conference on Electronic Commerce (EC-99)*, pages 148–157, 1999.
- [17] Kevin Roberts. The characterization of implementable choise rules. In Jean-Jacques Laffont, editor, *Aggregation and Revelation of Preferences*, pages 321–349. North-Holland, 1979. Papers presented at the first European Summer Workshop of the Econometric Society.
- [18] Amir Ronen. Mechanism design with incomplete languages. <http://task.stanford.edu/>, 2001. (full version).
- [19] Tuomas Sandholm. Issues in computational vickrey auctions. *International Journal of Electronic Commerce*, 4(3):107–129, 2000.
- [20] Tuomas W. Sandholm. Approaches to winner determination in combinatorial auctions. *Decision Support Systems*, to appear.

- [21] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.

## Appendix H

# Anonymous Bidding and Revenue Maximization

Yossi Feinberg\*and Moshe Tennenholtz†‡  
Stanford University

March 19, 2002

## 1 Introduction

The recent adoption of market mechanisms in general and auctions in particular for electronic commerce raises potent new theoretical questions. One aspect of an online environment is the prevalence of anonymous interaction. In particular, it is easier to maintain anonymity in an online auction than it is to maintain it in an offline auction. In the case of a single seller with multiple bidders anonymity mostly pertains to the identity of the bidders.

A natural question would be, should the seller opt for minimizing the opportunity for anonymous bidding? or, more precisely, when should the seller adopt an anonymous auction mechanism as a function of the interdependencies between the bidders' valuations? In this paper we show that even in the single unit english auction case, there seems to be no simple qualitative property that characterizes whether anonymous bidding yields higher or lower expected revenue to the seller.

The notion of anonymous bids employed here requires an additional explanation. We consider two variants of a dynamic ascending bid auction –

---

\*Stanford Graduate School of Business, Stanford CA 94305-5015. e-mail:  
yossi@gsb.stanford.edu

†Computer Science Department, Stanford University, Stanford CA 94305. e-mail:  
moshe@robotics.stanford.edu

‡This author gratefully acknowledges financial support by Darpa grant F30602-98-C-0214-P00005.

english auction. In the first mechanism the bidders actions are observable and each bidder can observe the identity of a bidder that dropped from the race. The second mechanism differs only in one aspect, the bidders only observe the fact that someone dropped from bidding at a given price, but they do not know who that bidder is. Admittedly, we consider a narrow view of the notion of anonymity and in particular exclude the discussion of uncertainty as to the identity of the players that *participate* in this auction.

Our example consists of three risk neutral bidders and a risk neutral seller. Two bidders, Ann and Bob, have independent identically distributed (uniform) private valuations and the third bidder, Carol, has a valuation equal to Bob's valuation plus a positive constant. Both Ann and Bob know their own valuation and the constant determining Carol's valuation as a function of Bob's valuation is commonly known. However, Carol does not know Bob's valuation ex-ante, i.e., does not know her own valuation. It turns out that both variants of the english auction support a unique perfect equilibrium. We compare the expected revenue for the seller at this equilibrium as a function of the constant determining Carol's valuation given Bob's valuation. The main result is that for some values of this constant the expected revenue is higher when the bidders observe the identity of a bidder that drops, while for other values it is the auction where bidders *do not* observe the identity of a dropping bidder that yields the higher expected revenue. The puzzling feature of this example is that the information structure and correlation structure are basically the same for every value of this variable. It is a quantitative change that determines which mechanism is more profitable to the seller rather than a qualitative one.

## 2 An Example

Consider 3 buyers Ann, Bob and Carol bidding for a single indivisible good in an ascending bid auction. Assume that the price  $p$  ascends from 0 to 1. Let  $v_A$  be Ann's valuation uniformly distributed in the interval  $[0, 1]$ . Let  $v_B$  be Bob's valuation which is independent of Ann's valuation and identically distributed. Both Ann and Bob know their own valuation. Let  $v_C$  be Carol's valuation which is equal to  $v_B + \alpha$  for some commonly known positive  $\alpha \in (0, 1)$ . Assume that Carol does not know her valuation. These distributions and the information available to the bidders are assumed to be commonly

known.

We consider two procedures for the ascending auction. In the first version each bidder can observe the identity of the other bidders, i.e., a bidder can identify who drops at a certain price. In the second procedure each bidder can only observe that someone dropped but the identity of the bidder that dropped is not revealed – the anonymous case.

**Lemma 1** *There exists a unique perfect equilibrium in pure strategies for each of the two procedures.*

We prove this lemma by explicitly constructing a pure strategy equilibrium for this example.

Let Ann's strategy be: "drop at  $p$  iff  $p > v_A$ "

Let Bob's strategy be "drop at  $p$  iff  $p > v_B$ "

Let Carol's strategy be "drop at  $p$  iff  $p > E^p$ " where  $E^p$  is the expectation of  $v_C$  given that Carol wins the auction at  $p$  and that Ann and Bob follow the strategies above. Note that  $E^p - p$  is actually Carol's expected payoff if she wins the auction at  $p$  and Ann and Bob follow the strategies above.

Since Ann and Bob are both perfectly informed as to their private valuation of the item, the strategies described above are weakly dominant strategies for them. Moreover Ann's strategy strictly dominates any other strategy at any given price  $p < v_A$ . We also note that Bob's strategy is strictly dominant when Carol does not observe the identity of a bidder that drops as long as  $p < v_B$ , and it is dominant whenever we perturb the other bidders' strategies. Hence both Ann and Bob play the unique perfect equilibrium strategies under the assumption that they play optimally at every price  $p$ . By definition, Carol's strategy is a best response to strategies ascribed to Ann and Bob. Thus we have the unique perfect equilibrium.  $\square$

We now explicitly calculate Carol's strategy for an arbitrary  $\alpha$ .

Consider the first case where Carol (and everyone else) can observe the identity of a bidder who drops from the auction. Recall that Carol's valuation is Bob's valuation plus  $\alpha$ . Hence she would bid as long as  $p \leq p_B + \alpha$  where  $p_B$  is the price where Bob dropped, i.e., it is equal to  $v_B$ . For each of these  $p$ 's her expected payoff *if she wins* is non-negative (it is strictly positive if the strict inequality holds). For every  $p > p_B + \alpha$  her expected payoff is strictly negative *if she wins*. But in this case  $E^p = p_B + \alpha$ . So Carol's strategy is (not surprisingly) to bid until the price is increased by  $\alpha$  from the point where Bob dropped.

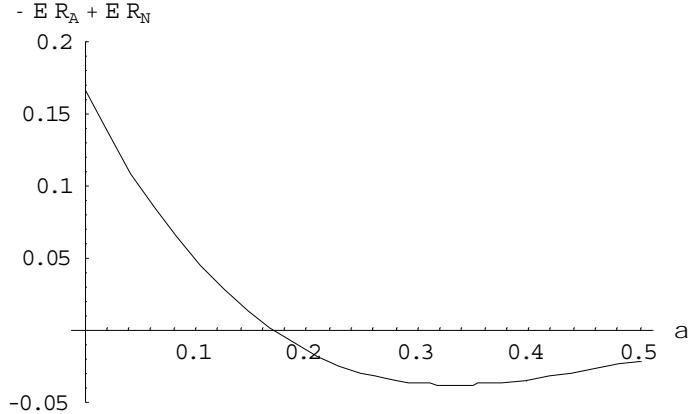


Figure 1:  $E(R_N) - E(R_A)$  as a function of  $\alpha$

In the second case Carol clearly stays as long as no one else drops. The moment one other bidder drops, say at the price  $p_1$ , Carol must assign a probability of .5 to  $v_B = p_1$  – the case it was Bob who dropped – and the rest of the weight is uniformly distributed on the interval  $[p_1, 1]$  – the case it is Ann who dropped. At every price  $p \geq p_1$  we have that Carol's belief as to the distribution of  $v_B$  has an atom at  $p_1$  with probability .5 and the rest of the weight is uniformly distributed on the interval  $[p, 1]$ . Obviously if another bidder drops then the auction is over. But the moment she wins the auction, say at the price  $p$ , her belief as to Bob's valuation is  $p_1$  with probability .5 and  $p$  with probability .5. Hence, her expected payoff is  $.5(p_1 + p) + \alpha - p = .5(p_1 - p) + \alpha$  and she will drop iff this payoff is negative. We just deduced that Carol will only drop the auction at  $p > p_1 + 2\alpha$ .

We now turn to the calculation of the seller's revenue as a function of  $\alpha$ .

In the non-anonymous case we have that  $R_N = \text{Max} \{ \text{Min} \{ v_A, v_B + \alpha \}, v_B \}$  according to the strategies above, and for the anonymous case we have  $R_A = \text{Min} \{ \text{Max} \{ v_A, v_B \}, \text{Min} \{ v_A, v_B \} + 2\alpha \}$ . The expected revenue to the seller is therefore  $E(R_N) = 1/2 + 1/2a - a^2/2 - a^3/6$  and  $E(R_A) = 1/3 + 2a - 4a^2 + 8a^3/3$  respectively. The graph depicted in Figure 1 plots  $E(R_N) - E(R_A)$  as a function of  $\alpha$ .

As claimed, for the given information structure, the anonymous mechanism yields a higher expected revenue for the seller for some values of  $\alpha$  (approximately higher than .171) and it yields a lower expected revenue for

other values of  $\alpha$  (below .171).

### 3 Discussion

One needs to be precise as to the sense in which the two mechanisms have similar information structures. For a given  $\alpha$  we actually use the same extensive form game with an interim refinement of the information structure as someone drops from the bidding process. The important feature is that the ex ante information structure is identical for both games. When varying  $\alpha$  we maintain the same game form for both mechanisms but change the payoffs in an identical manner for both the anonymous and the non anonymous auctions. One can also view this comparison as analyzing a single mechanism with a refined information structure. The characterizing feature of this refinement stems from the natural structure of an auction – the ability to observe the identity of a bidder that drops. It is interesting to note that even if we normalize (divide by the expected revenue) the difference between the expected revenue for the two mechanisms as a function of  $\alpha$ , we find that the normalized difference is not monotonic.

### References

## Appendix I

# From Belief Revision to Belief Fusion

Pedrito Maynard-Reid II and Yoav Shoham

Computer Science Department

Stanford University

Stanford, CA 94305

{pedmayn,shoham}@cs.stanford.edu

February 27, 2000

### Abstract

We introduce a new operator—*belief fusion*—which is a generalization of the classical AGM revision operator to the multi-agent case. In the process we define *pedigreed belief states*, which enrich standard belief states with the source of each piece of information. We show that AGM revision can be derived from belief fusion. We then note that the fusion operator defines a semi-lattice, and in particular is idempotent, associative, and commutative. As one consequence, we illustrate how belief fusion can be iterated without difficulty, in contrast to belief revision whose iteration has proved challenging. Finally, we define belief *diffusion*; whereas fusion produces a belief state with more information than is possessed by either of its two arguments, diffusion produces a state with less information. Fusion and diffusion are symmetric operators, and together define a distributive lattice.

## 1 Introduction

In what is by now classical work, Alchourrón, Gärdenfors, and Makinson [13, 1] proposed a theory of “reasonable” belief revision, the *AGM theory* henceforth. The intention of the theory is to formalize an Occam’s-razor principle, ensuring that beliefs change only when forced to by new information. The most common way of presenting the AGM theory is through the famous AGM postulates, which impose restrictions that attempt to capture this principle precisely.<sup>1</sup>

---

<sup>1</sup>Although the discussion in this paper will be semantic rather than axiomatic, for completeness we include here the AGM postulates (as formulated in [17] for the finite propositional case). If  $K$  is a theory in some propositional language,  $p$  and  $r$  are sentences in that language, and  $\circ$  is a revision operator, then:

R1  $K \circ p$  implies  $p$

R2 If  $K \wedge p$  is satisfiable, then  $K \circ p = K \wedge p$

There has been much subsequent work in several disciplines, consisting mostly of complaints about and modifications of the AGM postulates and setting. The catalyst for much of this work in recent years has been the iterated case of belief revision (that is, revising previously-revised beliefs). The AGM postulates restrict an agent's beliefs after a single revision, but provide no assistance in determining what an agent ought to believe after a sequence of revisions. Work on extending AGM to the iterated case includes [7, 10, 11, 19, 22, 26], but it is fair to say that as of now the theory of iterated belief revision is not a settled matter; see [12] for discussion of some of the outstanding issues.

Our direct interest lies in multi-agent belief revision, that is, the situation in which an agent is informed by multiple other agents, and, more interestingly, when multiple agents inform each other. However, it turns out that this issue is inextricably bound to that of iterated belief revision. Not only do we view multi-agent revision as a sequence of revisions each of a single agent's beliefs, but we will show that under the multi-agent perspective iterated belief change is unproblematic.

The basis for this paper is two observations, both of which are discussed further in the next section:

1. The AGM revision operator contains two asymmetries in its two arguments. The obvious asymmetry is the precedence of the second argument over the first one. The more subtle asymmetry, which is exposed only by examining the model theoretic characterization of the AGM setting, is the richer structure of the first argument as compared to the second.
2. The very setting of AGM revision is open to many interpretations, and resolving problems associated with AGM revision requires in general choosing among these interpretations. In particular, there is a choice between a *temporal perspective* and a *multi-agent perspective*.

We will adopt the multi-agent perspective, and will develop a theory of *belief fusion* which removes the second source of asymmetry from belief revision (but not, in this paper, the first asymmetry). Some of the specific contributions of this paper are as follows:

- The new fusion operator is technically and conceptually clear.
- Its definition appeals to another novel definition, of *pedigreed belief state*, which enriches the standard notion of belief state with the source of each belief.
- AGM revision can be derived from belief fusion.

---

R3 If  $p$  is satisfiable, then  $K \circ p$  is satisfiable

R4 If  $K_1 \equiv K_2$  and  $p_1 \equiv p_2$ , then  $K_1 \circ p_1 \equiv K_2 \circ p_2$

R5  $(K \circ p) \wedge r$  implies  $K \circ (p \wedge r)$

R6 If  $(K \circ p) \wedge r$  is satisfiable, then  $K \circ (p \wedge r)$  implies  $(K \circ p) \wedge r$

- Iterated fusion is not only well defined but also extremely well-behaved. This is because the fusion operator defines a *semi-lattice*, and in particular is idempotent, associative, and commutative.
- An additional operator is defined, *diffusion*, which is symmetric to fusion; whereas fusion in general adds information, diffusion removes some. Together, the fusion and diffusion operators define a *distributive lattice*.

We now proceed to cover these points in order.

## 2 The two asymmetries in AGM revision

Since the conceptual elements of our approach are as important as the technical ones, in this section we start with a somewhat lengthy pre-formal discussion of background and intuition. The remaining sections are mostly formal.

As classically presented, an AGM revision operator  $\circ$  accepts two arguments—a (typically, propositional logic) theory  $K$  and a sentence  $p$  in some language  $\mathcal{L}$ —and produces a new theory  $K \circ p$ . Or, from the semantic point of view, a revision operator is usually viewed as accepting two sets of interpretations—those satisfying  $K$  and  $p$ , respectively—and producing a third set, one satisfying  $K \circ p$ . As we shall discuss, this is a misleading view which is exposed by looking more closely at the semantics of AGM revision.

Indeed, the entire discussion in this paper will be semantic rather than axiomatic, and so it will be useful to start by recalling the well-known model theoretic characterization of AGM revision [14, 17]. Let  $\mathcal{W}$  be the set of worlds (i.e., interpretations) for  $\mathcal{L}$ . A revision operator  $\circ$  satisfies the AGM postulates if and only if for every theory  $K$  there exists a total pre-ordering  $\preceq$  over  $\mathcal{W}$  such that the worlds minimal with respect to  $\preceq$  are exactly those that satisfy  $K$  and, for every sentence  $p$ , the worlds that satisfy  $K \circ p$  are precisely the minimal worlds, with respect to  $\preceq$ , satisfying  $p$ . Indeed, the role of orderings in belief revision and non-monotonic logics has been well established in the literature.

In the sequel, we will call a pair  $(\mathcal{W}, \preceq)$  a belief *state*, and a set of worlds  $W \in \mathcal{W}$  a belief *set*. Intuitively, a belief set describes an agent’s actual beliefs, while a belief state describes his conditional belief sets given any possible new information. Clearly, every belief state induces a belief set, namely the set of minimal worlds in the belief state.

Although for those familiar with the AGM postulates the model theoretic characterization was obvious in hindsight, it has far-reaching ramifications. In particular, it means that *revision is a uniquely defined operation that takes as its first argument not a mere belief set, but a full belief state*. The AGM postulates are not rendered meaningless by this observation, but it is important to realize that they employ a misleading notational economy by implicitly building into the revision operator information more accurately considered as part of its first

argument. Consider, for example, the first postulate (see footnote in Introduction), ‘R1.  $K \circ p$  implies  $p$ .’ The naive reading of this postulate is “When the belief set validating  $K$  is revised by …”; the correct reading should be “When any belief *state* whose induced belief *set* validates  $K$  is revised by ....”<sup>2</sup> In the remainder of the paper, we assume “AGM setting” and “AGM revision” refer to this modified view, explicitly indicating any reference to the original view.

From this perspective, it is clear that the AGM setting contains two sources of asymmetry. First, as is well-known, the second argument to the revision operator takes complete precedence over the first one (see postulate R1 above). Second, as we have just discussed, the first argument is a full belief state, whereas the second is a mere belief set. This second asymmetry is more subtle and, we believe, ultimately deeper.

Some recent work in the area has attacked the first source of asymmetry. This asymmetry is often interpreted as “new information overrides old information,” and there have been suggestions that this chronological precedence is unjustified in general (recalling similar conclusions in the case of non-monotonic temporal reasoning, cf. [24]). However, it’s important to realize that there is nothing in the AGM setting to uniquely sanction the temporal interpretation. In particular, several researchers choose to view the process as one in which the belief sets of two agents are combined to produce a third. In this view, the first asymmetry amounts to giving one agent (the ‘expert’) total precedence over the other (the ‘novice’), and these recent attempts have been geared towards capturing less biased kinds of belief pooling. For example, [20] use the term *arbitration* to describe a commutative revision operator. In their system the fairness is achieved by omitting the first AGM axiom (R1 above) (they also consider adding other restrictions on arbitration, but these are not geared towards fairness). [18] place an additional fairness requirement that amounts to requiring that when two inconsistent theories are merged each one has to give up something. Other research in the area includes [3, 6, 23].

Since we agree with [12] that the AGM setting is unclear on issues of interpretation, we consider it meaningless to argue that one interpretation—the temporal one or the multi-agent one—is right and another wrong, only that one should be clear on one’s interpretation and should explore its consequences. However, we do argue that the multi-agent perspective leads to quite attractive properties.

We replace the operator of belief revision by the operator of *belief fusion*. Like merging and arbitration, fusion involves two agents, whose beliefs are fused. Specifically, the arguments to belief fusion are two full belief states. Unlike merging and arbitration, however, there is nothing fair about fusion. Indeed, in a precise sense fusion is a faithful generalization of AGM revision to the multi-

---

<sup>2</sup>One important change is necessary: We rewrite R4 as “ $\Psi_1 = \Psi_2$  and  $p_1 \equiv p_2$ , then  $K_1 \circ p_1 \equiv K_2 \circ p_2$ ,” where  $\Psi_1$  and  $\Psi_2$  are belief states. Without this change which, essentially, allows the result of a revision to depend on past revisions, most iterated revision proposals—including our own—are inconsistent with the postulates.

agent case. We show that fusion is extremely well-behaved, and in particular, its iteration poses no problems.

Since in this paper we do not directly challenge the first asymmetry and continue to rely on dominance in the process of fusion, two remarks are in order. First, as we shall see, the notion of dominance we have is much more fine-grained than that of one agent dominating another. Essentially, we will have one agent dominating another only with respect to particular judgments. Second, while our framework can be adapted to embrace ideas on “fair” merging, it is instructive to see that solving the problems that have plagued iterated belief revision does not require doing away with dominance as a method for resolving conflicting beliefs.

A proposal related to our own is that of [8] for revising belief states by conditional beliefs. That work can be thought of as taking into account not necessarily the revising agent’s unconditional belief, but his conditional ones. In a sense our construction takes into account his entire set of conditional beliefs. Other differences include the fact that our approach also takes into consideration sources of information and the relative credibility of these sources. Finally, we think it a fair statement that our approach is based on clearer semantical underpinnings.

Perhaps closest to our work is the recent proposal in [9] for combining information from conflicting sources. He addresses a complementary problem to our own: deciding what information to reject given the subset of informing sources rejecting the information. In making this decision, Cantwell assumes a generalization of our credibility ordering, in this case a partial pre-order over sets of sources. He explores a number of ways of inducing a partial pre-order over sentences based on this ordering, which can then be used to determine a subset (although not all) of the sentences to reject. The proposal also differs with ours in that the sources of information and resulting belief states are essentially belief sets; non-trivial conditional beliefs are not accounted for. In addition, the work does not address the problem of combining these belief states. The degree to which the framework captures our intuitions in specific domains deserves further research.

Other related research include [4] which approaches information aggregation from a possibilistic logic point of view, and several papers in a special issue of *Theoria* [15] which also seek to extend the AGM framework to deal with non-prioritized revision.

### 3 Belief fusion

First, a bit of standard notation: We assume some language  $\mathcal{L}$ . A world  $w$  is an interpretation over  $\mathcal{L}$ , and we say that for a sentence  $p \in \mathcal{L}$ ,  $w \models p$  iff  $p$  evaluates to true in  $w$ . Given a set of worlds  $\mathcal{W}$  and a sentence  $p$ ,  $\|p\| = \{w \in \mathcal{W} \mid w \models p\}$ . If  $p$  and  $r$  are sentences, then  $p \models r$  iff  $\forall w \in \|p\|, w \models r$ . Also, in the treatment

that follows we make use of a number of (pre-)orders. Given a set  $\Theta$  and a (pre-)order  $\leq$  over  $\Theta$ , we define  $\min(\leq, \Theta) = \{\alpha \in \Theta \mid \forall \beta \in \Theta, \alpha \leq \beta\}$ .

Let us start by formally defining belief states. For reasons that will be made clear, we shall call them *anonymous* belief states.

**Definition 1** An (anonymous) belief state (over  $\mathcal{W}$ ) is a pair  $(\mathcal{W}, \leq)$  where  $\mathcal{W}$  is a set of possible worlds and  $\leq$  is a total pre-order over  $\mathcal{W}$ .

We use  $<$  to denote the strict version of  $\leq$ . In this article the set  $\mathcal{W}$  will not play a role, and can be assumed to be fixed. We denote by  $s_0$  the ‘agnostic’ belief state, in which  $\leq$  is the complete relation.

To first approximation, the belief fusion operator we will define accepts two belief states and produces a third. However, in order for the operator to be meaningful, it will require additional input, which, intuitively, will adjudicate between the two belief states where they disagree.

It is tempting to resolve conflicts by declaring one agent ( $B$ ) more credible than the other ( $A$ ) and have his judgments dominate. Specifically, one could define fused belief state  $A \circledast B$  to be the *refinement* of  $B$  by  $A$ . Here is the definition of this straw-man fusion operator,  $\circledast$ :

$$A \circledast B = \{(w_1, w_2) : (w_2, w_1) \notin B \vee ((w_1, w_2) \in B \wedge (w_2, w_1) \notin A)\}.$$

In other words, we would construct the fused belief state as follows: for each pair of worlds, whenever the more credible agent strictly prefers one world to the other, we side with this preference. In cases where the most credible agent has no preference, we follow the ranking of the less credible agent. Naturally,  $\circledast$  is not a symmetric operator. This operator is illustrated in Figure 1. The dots labeled with lower-case letters are worlds; the circles represent equivalence classes of worlds.

Figure 1: The straw-man fusion operator (belief sets in each belief state are highlighted).

This is a well-defined operation in that it produces a total pre-order. However, there is a problem with this definition pertaining to the iteration of the

operator. Consider three belief states A, B and C with increasing order of dominance (A dominated by B, and both by C). Presumably, the above definition would give meaningful interpretation to  $(A \otimes^* B) \otimes^* C$ , since, intuitively speaking, all the information in C dominates all the information in  $A \otimes^* B$ . But what about  $(A \otimes^* C) \otimes^* B$ ? Here it would seem that some of the information in  $A \otimes^* C$  dominates the information in B (because it originated from C) and some is dominated by it (because it originated from A).

The problem is that the standard belief state is not rich enough to represent the source of each information item, which is the reason we term it ‘anonymous’. Our actual definition will enrich belief states with this missing information. To develop intuition for the following definitions, imagine a set of information sources and a set of agents. The sources can be thought of as primitive agents with fixed (anonymous) belief states. Each source informs some of the agents of its belief state; in effect, each source offers the opinion that certain worlds are more likely than others, and remains neutral about other pairs.

An agent’s belief state is simply the amalgamation of all these opinions, each annotated by its origin (or “pedigree”). Of course, these opinions in general conflict with one another, and the agent must resolve these conflicts in order to arrive at a coherent belief state. There are various plausible ways of performing this resolution. In this paper we assume that the agent places a strict “credibility” ranking on the sources, and accepts the highest-ranked opinion offered on every pair of worlds.

The following definition considers only finite sets of sources; this restriction can be relaxed at the price of complicating the subsequent development in this paper.

**Definition 2** *Given a finite set of anonymous belief states  $S \subset \mathcal{S}$  the pedigree belief state (over  $\mathcal{W}$ ) induced by  $S$  is a function  $\Psi : \mathcal{W} \times \mathcal{W} \mapsto 2^{S \cup \{s_0\}}$  such that*

$$\Psi(w_1, w_2) = \{(\mathcal{W}, \leq) \in S : w_1 < w_2\} \cup \{s_0\}.$$

We will use  $\mathcal{S}$  to denote the set of all of sources over  $\mathcal{W}$ , and throughout this paper we will consider pedigree belief states that are induced by subsets of  $\mathcal{S}$ . Note that both  $\{\}$  and  $s_0$  induce the same pedigree belief state; we will denote it too by  $a_0$ . Finally, we will use  $a_{\max}$  to denote the pedigree belief state induced by  $\mathcal{S}$ .

Next we define a particular policy for resolving conflicts within a pedigree belief state. We assume a strict ranking  $\sqsubseteq$  on  $\mathcal{S}$  (and thus also on the sources that induce any particular  $\Psi$ ); the strictness of the ranking is a significant restriction that we discuss further in the final section. We interpret  $s_1 \sqsubseteq s_2$  as ‘ $s_2$  is more credible than  $s_1$ ’. As usual, we define  $\sqsubseteq$ , read “as credible as”, as the reflexive closure of  $\sqsubseteq$ .

We also assume that  $s_0$  is the least credible source, which may merit some explanation. It might be asked why equate the most agnostic source with the

least credible one. In fact we don't have to, but since in the definitions that follow, agnosticism is overridden by any opinion regardless of credibility ranking, we might as well assume that all agnosticism originates from the least credible source, which will permit simpler definitions.

Intuitively, given a pedigreed belief state  $\Psi$ ,  $\Psi_{\sqsubset}$  will retain from  $\Psi$  the highest-ranked opinion about the relative likelihood between any two worlds.

**Definition 3** *Given  $\mathcal{W}, \mathcal{S}, \Psi$  and  $\sqsubset$  as above, the dominating belief state of  $\Psi$  is the function  $\Psi_{\sqsubset} : \mathcal{W} \times \mathcal{W} \mapsto S$  such that  $\forall w_1, w_2 \in \mathcal{W}$  the following holds: If  $\max(\Psi(w_2, w_1)) \sqsubset \max(\Psi(w_1, w_2))$  then  $\Psi_{\sqsubset}(w_1, w_2) = \max(\Psi(w_1, w_2))$ . Otherwise,  $\Psi_{\sqsubset}(w_1, w_2) = s_0$ .*<sup>3</sup>

Clearly, for any  $w_1, w_2 \in \mathcal{W}$  either  $\Psi_{\sqsubset}(w_1, w_2) = s_0$  or  $\Psi_{\sqsubset}(w_2, w_1) = s_0$  or both.

Somewhat surprisingly,  $\Psi_{\sqsubset}$  induces a standard (anonymous) belief state:

**Definition 4** *The ordering induced by  $\Psi_{\sqsubset}$  is the relation  $\preceq \subseteq \mathcal{W} \times \mathcal{W}$  such that  $w_1 \preceq w_2$  iff  $\Psi_{\sqsubset}(w_2, w_1) = s_0$ .*

We denote the strict version of  $\preceq$  by  $\prec$ .

**Proposition 1**  $\preceq$  is a total pre-order on  $\mathcal{W}$ .

Thus a dominating belief state is a generalization of the standard notion of (anonymous) belief state, representing the agent's ordering on worlds based on the agent's opinion of their relative likelihood as well as which source the opinion originated from. Now, if the agent later interacts with another agent and they disagree over some piece of information, intuitively they can resolve the conflict based on who has the stronger support.

The fusion operator we define captures this intuition. We first give a very natural definition for the fusion of two pedigreed belief states  $\Psi_1$  and  $\Psi_2$  based on their respective sets of supporting sources: we simply combine them. Then we show that it is possible to compute the new pedigreed belief state directly in terms of the  $\Psi_1$  and  $\Psi_2$  without needing to refer to the sets of sources. Furthermore, we show how to determine the new dominating belief state based on those associated with  $\Psi_1$  and  $\Psi_2$ . As it turns out, the result will match the conflict-resolution policy we outlined above.

**Definition 5** *Given a set of sources  $\mathcal{S}$  and  $\sqsubset$  as above,  $S_1, S_2 \subset \mathcal{S}$ , the pedigreed belief state  $\Psi_1$  induced by  $S_1$ , and pedigreed belief state  $\Psi_2$  induced by  $S_2$ , the fusion of  $\Psi_1$  and  $\Psi_2$ , denoted  $\Psi_1 \odot \Psi_2$ , is the pedigreed belief state induced by  $S_1 \cup S_2$ .*

---

<sup>3</sup>Note the use of the restrictions. Finiteness assures that a maximal source exists; we could readily replace it by weaker requirements on the infinite set. The absence of ties in the ranking  $\sqsubset$  ensures that the maximal source is unique; removing this restriction is not straightforward.

Obviously, the set of pedigreed belief states is closed under  $\otimes$

**Proposition 2**

1.  $(\Psi_1 \otimes \Psi_2)(w_1, w_2) = \Psi_1(w_1, w_2) \cup \Psi_2(w_1, w_2)$
2.  $(\Psi_1 \otimes \Psi_2)_{\sqsubset}(w_1, w_2) =$   
 $\max(\Psi_{1\sqsubset}(w_1, w_2), \Psi_{2\sqsubset}(w_1, w_2))$   
 $\text{if } \max(\Psi_{1\sqsubset}(w_2, w_1), \Psi_{2\sqsubset}(w_2, w_1)) \sqsubset$   
 $\max(\Psi_{1\sqsubset}(w_1, w_2), \Psi_{2\sqsubset}(w_1, w_2)), \text{ and}$   
 $s_0 \text{ otherwise.}$
3. If  $\preceq_1, \preceq_2, \text{ and } \preceq$  are the orderings induced by  $\Psi_{1\sqsubset}, \Psi_{2\sqsubset}$ , and  $(\Psi_1 \otimes \Psi_2)_{\sqsubset}$ , respectively, then  
 $w_1 \prec w_2 \text{ iff}$   
 $w_1 \prec_1 w_2 \text{ and } \max(\Psi_{2\sqsubset}(w_1, w_2), \Psi_{2\sqsubset}(w_2, w_1)) \sqsubseteq \Psi_{1\sqsubset}(w_1, w_2) \text{ or}$   
 $w_1 \prec_2 w_2 \text{ and } \max(\Psi_{1\sqsubset}(w_1, w_2), \Psi_{1\sqsubset}(w_2, w_1)) \sqsubseteq \Psi_{2\sqsubset}(w_1, w_2).$

The second property formalizes the idea that, for a given pair of worlds, the new dominating belief state should choose the order of the pair that gives the most credible support between the two agents for this pair of worlds, assigning the same support to this order, and  $s_0$  to the opposite order. The third property describes how to derive the new induced ordering from those of the two fused belief states.

Figure 2 illustrates the fusion operation on three dominating belief states. We can view  $A$ ,  $B$ , and  $C$ , to be agents with information of from sources of high (source 3), medium (source 2), and low (source 1) credibility, respectively. Because we now make precedence decisions at a local rather than global level based on sources of support, fusing  $A$  with  $C$  and the result with  $B$  is now well-defined in a conceptually justified way, unlike in the case of the strawman operator discussed earlier. Notice the dependence of the final belief state on all three sources.

We will further explore the properties of belief fusion in later sections, but first we discuss the connection between belief fusion and classical AGM revision.

## 4 Revision as under-specified fusion

Now that we have defined fusion, one can view the traditional AGM revision operator as the application of the fusion operator to a partially specified input (only the belief set of the expert is given, not his full belief state). In general, the full belief state of the expert strongly affects the resulting “fused” belief state. However, it turns out that the belief set defined by the fused belief state depends only on the belief set of the expert. We now show that this is so, and that the AGM revision precisely captures the properties of this belief set.

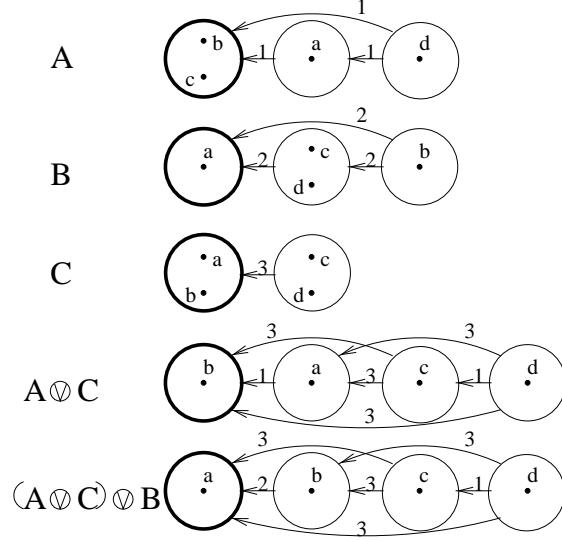


Figure 2: The correct fusion operator.

In order to mimic AGM revision we need to be able to differentiate between the expert agent and the novice. We do so by defining an ordering on agents, (or, equivalently, on pedigreed belief states). Intuitively, one can distinguish between the *quantity* of information an agent has (which worlds he can distinguish) and its *quality* (what are the sources of these distinctions). The following definition ranks agents first on quality, breaking ties by quantity:

**Definition 6** *Agent  $A_2$  with a pedigreed belief state  $\Psi_2$  over  $\mathcal{W}$  has as reliable sources as agent  $A_1$  with pedigreed belief state  $\Psi_1$  over  $\mathcal{W}$  (written  $A_2 \geq A_1$  or  $\Psi_2 \geq \Psi_1$ ) iff it is the case that whenever  $\Psi_{2\sqsubseteq}(w_1, w_2) \neq s_0$  then  $\max(\Psi_{1\sqsubseteq}(w_1, w_2), \Psi_{1\sqsubseteq}(w_2, w_1)) \sqsubseteq \Psi_{2\sqsubseteq}(w_1, w_2)$ .*

**Proposition 3** *Let  $\Psi_1$  and  $\Psi_2$  be pedigreed belief states, and let  $\preceq_1$  and  $\preceq_2$  be the orderings induced by  $\Psi_{1\sqsubseteq}$  and  $\Psi_{2\sqsubseteq}$ , respectively. Further, let  $\preceq$  be the ordering induced by  $(\Psi_1 \odot \Psi_2)\sqsubseteq$ . If  $\Psi_2 \geq \Psi_1$ , then  $w_1 \prec_2 w_2$  implies  $w_1 \prec w_2$  for all  $w_1, w_2 \in \mathcal{W}$ .*

Note that any pedigreed belief state has as reliable sources as  $a_0$ , and  $a_{\max}$  has as reliable sources as any other pedigreed belief state.

In the following, we use the notation  $\Psi\downarrow$  to denote the belief set defined by a pedigreed belief state  $\Psi$ , that is, the set of worlds minimal with respect to the ordering induced by  $\Psi\sqsubseteq$ . Also, we use  $\Psi \circ \omega$  to denote the revision of belief

state  $\Psi$  by belief set  $\omega$  according to AGM (that is, the worlds in  $\omega$  that are minimal according to  $\Psi_{\square}$ ).

**Proposition 4** *Let  $\Psi_1$  and  $\Psi_2$  be pedigreed belief states such that  $\Psi_2 \geq \Psi_1$ . Then  $(\Psi_1 \odot \Psi_2) \downarrow = \Psi_1 \circ (\Psi_2 \downarrow)$ .*

**Corollary 4.1** *Let  $\Psi_1, \Psi_2, \Psi_3$  be pedigreed belief states such that  $\Psi_2 \geq \Psi_1$ ,  $\Psi_3 \geq \Psi_1$ , and  $\Psi_2 \downarrow = \Psi_3 \downarrow$ . Then  $(\Psi_1 \odot \Psi_2) \downarrow = (\Psi_1 \odot \Psi_3) \downarrow$ .*

Thus, AGM revision is simply a projection of belief fusion, in which one ignores all but the belief set of one of the initial belief states, and all but the belief set of the resulting belief state.

## 5 Well-behavedness of iterated, multi-agent belief fusion

We mentioned in the introduction that the problem of iteration has proved a major challenge to AGM-style revision. We now show that this is not the case for fusion. To begin with, note that iteration is formally well-defined; the output of fusion (a pedigreed belief state) is a legitimate input to another fusion operation.

From the set-theoretic definition of fusion, it follows immediately that iterated belief fusion is not only well-defined, but also extremely well-behaved. In particular, it inherits the idempotence, commutativity, and associativity properties of  $\cup$ .

To demonstrate the well-behavedness of iterated belief fusion, we give several related examples which depend on these properties; the examples are stated informally for readability, but can easily be stated formally and proved. In all of them assume that there are  $n$  agents, each with his own belief state over the same set of worlds  $\mathcal{W}$ , all agreeing on their expertise ranking relative to one another, and all employing belief fusion as the method of update.

- One of the agents is the manager. Question: Will the order in which he gets briefed by his various employees affect his resulting belief state? Answer: No.
- The same manager is considering whether to get directly updated by the employees, or to have his vice-manager get updated by the rest of the employees, and then have the vice-manager update him. Question: Should the manager worry that the result will be skewed by the vice-manager's personal biases? Answer: No, the manager's resulting belief state will be as in the first case.

- The manager is gone and the team needs to reach consensus. Each agent broadcasts his belief state to all the others who receive it immediately, and incorporates all the belief states communicated to it. Question: Will the agents end up with identical belief states? Answer: Yes, even if they each perform the fusions in different orders.
- The situation is as above, but agents don't have unlimited broadcast capability. Instead, each agent can communicate with some of the others, and this capability is not necessarily symmetric. Question: Can the agents reach consensus through a process of fusion? Answer: Yes, iff the communication graph is strongly connected (the communication graph is the directed graph in which agents are nodes and directed arcs represent communication capability; a directed graph is strongly connected if there is a directed path from any node to any other). In this case each agent should simply communicate his belief state to all the agents he can, incorporate the belief states communicated to him, and repeat. After  $d$  rounds all agents will have identical belief states, where  $d$  is the diameter of the communication graph (the diameter of a directed graph is the longest shortest directed path between any two nodes in the graph).

## 5.1 Comparison to iterated revision approaches

It is natural to ask why we do not simply extend one of the existing iterated belief revision approaches to accommodate a multi-agent point of view. Specifically, we could assume that both arguments to an operator are full belief states, but that only the belief set portion of the second argument is used during revision. (Obviously, associativity does not make much sense given the temporal interpretation of iterated revision as the arguments are of different types, a belief state and a belief set.<sup>4</sup>) Accordingly, we briefly take a look at some of the recent iterated revision proposals, extended as described above, and subject them to one of the most benign invariance tests imaginable, namely, associativity.

We should point out that we don't necessarily view the invariance of associativity as obviously valid, even given a multi-agent interpretation; experience has taught us to be wary of postulates resting on loose intuition alone. But associativity is a natural criterion to consider, and it is interesting to see—even without attaching a value judgment to the outcome—how these proposals fare relative to this criterion.

---

<sup>4</sup>Incidentally, if we consider the original AGM postulates as applied to the revision of one theory by another, the question of whether associativity holds is legitimate. However, a simple example shows that associativity is actually inconsistent with the postulates: Consider the two possible associations of  $p$  revised by  $r$  revised by  $p \text{ XOR } r$ . If associativity is assumed, the AGM postulates—in particular, R1, R2, and R4—force the result of left association to entail  $\neg p \wedge r$ , and the result of right association to entail  $p \wedge \neg r$ , a contradiction. This can be traced to the independence of the original AGM revision on past revisions. However, the iterated revision operators we consider here are, like fusion, history-dependent.

As it so happens, not only does this invariance not hold for any of the proposals, but in each case it is possible to even get conflicting results depending on the revision order. We consider here the proposals in [7, 10, 25, 19, 26]. We describe each of these proposals below and show that there exists at least one example such that (a) all five proposals agree on the result of iterated revision, for any fixed association order of revision, and (b) these different orders of revision yield belief sets that are not only distinct, but actually mutually inconsistent. This counter-example is shown in Figure 3.

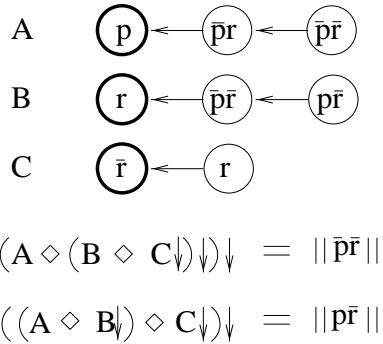


Figure 3: Counter-example showing alternative iterated revision operators are not associative.  $A, B, C$  are belief states.

**Boutilier's natural revision** *Natural revision*, proposed by Boutilier [7], extends the AGM idea of minimally changing beliefs to apply to the agent's counterfactual beliefs as well. Given a belief state, this approach specifies that we only change the ordering as much as is required by the AGM postulates, and no more.

**Proposition 5** *The resulting belief sets using left and right association of Boutilier's natural revision operators can be inconsistent.*

**Darwiche and Pearls' formulation** Darwiche and Pearl [10] suggest additional postulates in an attempt to adapt the AGM framework for iterated revision. As in the case of natural revision, this approach derives its inspiration from a notion of minimizing change to the belief state. However, it relaxes the constraint that change must be completely minimized, thus allowing for a whole family of revision operators, including natural revision as one instantiation. Somewhat surprisingly, natural revision's lack of associativity applies to every member in the family of operators.

**Proposition 6** *The resulting belief sets using left and right association of any revision operators satisfying the Darwiche and Pearl postulates can be inconsistent.*

**Spohn's conditionalization** In [25], Spohn introduces conditionalization operators over *ordinal conditional functions* (OCFs). OCFs can be viewed as anonymous belief states imbued with the additional structure of an ordinal ranking (aka a  $\kappa$ -ranking) over worlds so that it is possible to speak about degrees of belief.<sup>5</sup> Spohn proposed conditionalization operators over these functions as qualitative versions of probabilistic conditionalization. The set of operators is parameterized by  $\alpha$  which takes on ordinal values. Intuitively, revising by a sentence  $p$  using a particular  $\alpha$ -conditionalization operator will cause the agent to believe  $p$  with  $\alpha$  firmness.

**Proposition 7** *The resulting belief sets using left and right association of any combination of  $\alpha$ -conditionalization operators can be inconsistent.*

We hasten to point out that in this paper Spohn also defines an operator for the conditionalization of one OCF by another. With intuitions based on Jeffrey's generalized conditionalization [16], this operator is associative, though not commutative (given two OCFs  $\kappa$  and  $\lambda$ ,  $\kappa$  conditioned by  $\lambda$  generally is not the same as  $\lambda$  conditioned by  $\kappa$ ). The operator behaves quite similarly to ours in the special case where the conditioning agent's sources are all more reliable than the conditioned agent's.

**Lehmann's formulation** In [19], Lehmann proposes yet another set of postulates intended to regulate sequences of revisions. He provides a semantic account based on what he calls *widening rank models* which, like OCFs, can be viewed as augmented anonymous belief states. He provides a recursive definition for computing the result of a sequence of revisions based on a given widening rank model.

**Proposition 8** *The resulting belief sets using left and right association of Lehmann's revision operator can be inconsistent.*

**Williams' transmutations** Williams [26] generalizes Spohn's notion of conditionalization operators to include any operators over OCFs that satisfy the AGM properties, referring to this larger class of operators as the set of *transmutations*. She describes two particular sub-classes of transmutations: *conditionalization operators* which are equivalent to Spohn's conditionalization operators, and *adjustment operators*, a family of operators parameterized by  $\beta$  which takes

---

<sup>5</sup>Spohn actually defined OCFs with respect to subfields of  $2^{\mathcal{W}}$  closed under  $\cup$  and  $\cap$ . However, the additional structure does not play a role in our results and so, for the sake of clarity, we use the simpler definition.

on ordinal values. We have already seen that conditionalization operators are not associative. As it turns out, the same is true for adjustment operators.

**Proposition 9** *The resulting belief sets using left and right association of any combination of  $\beta$ -adjustment operators can be inconsistent.*

## 6 The belief lattice: fusion and diffusion

In the previous section we mentioned that  $\oslash$  is idempotent, commutative, and associative. Thus, a set of pedigreed belief states that is closed under  $\oslash$  forms a *semi-lattice* [5]. Intuitively, higher states in the lattice contain more information than lower ones (where, as explained, ‘more’ is determined first by quality and then by quantity).  $\oslash$  accepts two pedigreed belief states and returns the least pedigreed belief state that contains at least as much information as both of them. Note that this semi-lattice has a “unit” element,  $a_0$  (since  $\Psi \oslash a_0 = \Psi$ ) and an “annihilator” element,  $a_{\max}$  (since  $\Psi \oslash a_{\max} = a_{\max}$ ).

This suggests that there might be a symmetric operator to  $\oslash$  one which takes two pedigreed belief states and returns the greatest state containing no more information than either one. In fact, this operator can be readily defined:

**Definition 7** *Given  $S_1, S_2 \subset \mathcal{S}$ , the pedigreed belief state  $\Psi_1$  induced by  $S_1$ , and pedigreed belief state  $\Psi_2$  induced by  $S_2$ , the diffusion of  $\Psi_1$  and  $\Psi_2$ , denoted  $\Psi_1 \oslash\Psi_2$ , is the pedigreed belief state induced by  $S_1 \cap S_2$ .*

In other words, we transform fusion into diffusion by replacing the union of the sources by their intersection.

Trivially, we have the characterization of  $\Psi_1 \oslash\Psi_2$  directly in terms of  $\Psi_1$  and  $\Psi_2$ .

**Proposition 10**

$$(\Psi_1 \oslash\Psi_2)(w_1, w_2) = \Psi_1(w_1, w_2) \cap \Psi_2(w_1, w_2).$$

However, unlike the case of fusion, it is not possible to provide a characterization of  $(\Psi_1 \oslash\Psi_2)_{\sqsubseteq}$  (or its induced orderings) directly in terms of  $\Psi_{1\sqsubseteq}$  and  $\Psi_{2\sqsubseteq}$  (or their induced orderings); the latter simply do not contain enough information. This is illustrated in Figure 4. The figure shows the diffusion of two pedigreed belief states along with the corresponding dominating belief states. Now, consider the case where agent  $B$  also had source 1 as a source, i.e.,  $S_B = \{1, 2, 3\}$ . Although the dominating belief states for  $A$  and  $B$  would be identical to those in the figure, the dominating belief state resulting from diffusion would be exactly that of  $A$ . Thus, it is impossible, in general, to determine the new diffused state given solely the dominating belief states.

Clearly,  $\oslash$  also forms a semi-lattice. However, the roles of  $a_0$  and  $a_{\max}$  are reversed: the “unit” element is  $a_{\max}$  ( $\Psi \oslash a_{\max} = \Psi$ ) and the “annihilator” element is  $a_0$  ( $\Psi \oslash a_0 = a_0$ ). Also note that, together, the fusion and diffusion

operators represent a *distributive lattice* [5] over the set of pedigreed belief states. In particular, they are absorbitive and distributive.

## 7 Future work

We summarized the main contributions of this paper in the introduction, and discussed related research in the first two sections. We conclude here with a discussion of several of the many directions in which this work can be extended.

The restriction precluding equally ranked sources is an important one. Its root resides in the fact that it is unclear what to do in situations where sources of equal credibility offer conflicting information. One possibility would be to take the disagreement as reason for agnosticism. However, such a policy can lead to a loss of transitivity so that the result of a fusion is no longer another dominating belief state. That technicality aside, the issue resurfaces if later we are informed by a lower ranked source that has definite opinions on the matter. One could override the agnosticism as we have in the treatment above, thereby essentially promoting the opinion of the lower ranked source over the combined opinion of the higher ranked ones. These may seem reasonable—we might consider the less credible source to be a tie-breaker. However, the approach breaks down if instead of having two equally-ranked sources with opposite opinions, we had one hundred that voted one way and one that voted the other way. This would also result in a tie. If a lower ranked source came along later and sided with the one renegade source, the fusion operator would force agreement with it.

One could, of course, invent more clever schemes such as voting with the majority or using the next highest opinionated sources to break deadlocks. However, without weakening some basic assumptions, these will all be doomed in general, since it is possible to view our setting as a generalization of the setting Arrow addressed in his Impossibility Theorem [2]. Basically, we can model his setting as one where all agents are informed by  $n$  equally-ranked sources.<sup>6</sup> We are essentially asking that the following conditions hold:

- unrestricted domain: sources can be arbitrary total pre-orders over  $\mathcal{W}$ ,
- restricted range: the belief state induced by a set of these sources should be another total pre-order,
- independence of irrelevant alternatives: the ordering between two worlds in an induced belief state should only depend on how the sources rank those two worlds,
- weak Pareto principle: if all sources strictly prefer one world to another, this preference should be preserved, and

---

<sup>6</sup>More accurately, in his formulation, each agent is informed by  $n$  “individuals” where each individual’s belief state can be any of the possible sources. The distinction is not important here, however.

- nondictatorship: since the sources are equally-ranked, no particular source should have its opinions dominate.<sup>7</sup>

Arrow proved that there is no policy that obeys all of these conditions.

It is clear, however, that since *pedigreed* belief states retain the full pedigree of each belief, it is possible to experiment with many kinds of induced beliefs states other than dominant ones. In particular, in the second section we discussed recent interest in “fair” merging of beliefs. It will be interesting to see if we can capture the specific proposals made recently, and if not why.

In this paper when we assumed that all agents share the credibility ranking on sources. In general, and these rankings can vary among agents, and even change within an agent over time. Furthermore, an agent’s ranking function can depend on the context; different sources may have different areas of expertise. Exploring the behavior of fusion and diffusion in these more general settings is an obvious next step.

The work here has been qualitative in nature. However, often domains of interest have additional quantitative structure (e.g., a probability distribution rather than a simple total pre-order over worlds defining a belief state) which agents can take advantage of when modifying their mental states. Consequently, extending this work to provide principled accounts of how the belief states of a group of agents change under such conditions is another important followup step.

Finally, we note that while through this paper we viewed the (pre- or strict) orderings on possible worlds as describing ‘belief’, in fact there is nothing in the formalism to make the ‘preference’ interpretation less apt (indeed, this remark applies to most of the work in AI on belief revision and nonmonotonic reasoning). This raises the question whether there is an interesting connection to be made between the development in this paper and classical work in economics on preference aggregation.

## Acknowledgments

We would like to thank Daniel Lehmann and the members of the Nobots research group at Stanford University for insightful feedback. Pedrito Maynard-Reid II’s work was partly supported by a National Physical Science Consortium Fellowship. Yoav Shoham’s work was partly supported by the National Science Foundation under Grant No. IIS-9714701.

---

<sup>7</sup>This is actually a stronger restriction than that made by Arrow. Given his formulation in terms of individual variables that take on sources as values, the condition states that no one individual should always dominate. Here we are demanding that no individual *ever* dominate.

## References

- [1] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [2] Kenneth J. Arrow. *Social Choice and Individual Values*. Wiley, New York, 2nd edition, 1963.
- [3] C. Baral, S. Kraus, J. Minker, and V. S. Subrahmanian. Combining knowledge bases consisting of first-order theories. *Computational Intelligence*, 8(1):45–71, 1992.
- [4] S. Benferhat, D. Dubois, and H. Prade. From semantic to syntactic approaches to information combination in possibilistic logic. In B. Bouchon-Meunier, editor, *Aggregation and Fusion of Imperfect Information, Studies in Fuzziness and Soft Computing*, pages 141–161. Physica Verlag, 1997.
- [5] Garrett Birkhoff. *Lattice Theory*, volume 25 of *Colloquium Publications*. American Mathematical Society, 1948.
- [6] A. Borgida and T. Imielinski. Decision making in committees: A framework for dealing with inconsistency and non-monotonicity. In *Proceedings of the Workshop on Nonmonotonic Reasoning*, pages 21–32, 1984.
- [7] Craig Boutilier. Iterated revision and minimal change of conditional beliefs. *Journal of Philosophical Logic*, 25:263–305, 1996.
- [8] Craig Boutilier and Moisés Goldszmidt. On the revision of conditional belief sets. In G. Crocco, L. Fariñas del Cerro, and A. Herzig, editors, *Conditionals: From Philosophy to Computer Science*, chapter 9, pages 267–300. Oxford University Press, 1995.
- [9] John Cantwell. Resolving conflicting information. *Journal of Logic, Language, and Information*, 7:191–220, 1998.
- [10] Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89(1–2):1–29, 1997.
- [11] Michael Freund and Daniel Lehmann. Belief revision and rational inference. Technical Report TR 94-16, Hebrew University, 1994.
- [12] Nir Friedman and Joseph Y. Halpern. Belief revision: A critique. In *Proceedings of the Fifth International Conference on Principles of Knowledge Representation and Reasoning (KR '96)*, pages 421–431, 1996.
- [13] Peter Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, 1988.

- [14] Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.
- [15] Sven Ove Hansson. *Theoria: A Swedish Journal of Philosophy*, 63(1–2), 1997.
- [16] R. C. Jeffrey. *The Logic of Decision*. University Press, Chicago, 1965.
- [17] Hirofumi Katsuno and Alberto O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3):263–294, 1991.
- [18] Sébastien Konieczny and Ramón Pino Pérez. On the logic of merging. In *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR '98)*, pages 488–498, 1998.
- [19] Daniel Lehmann. Belief revision, revised. In *Proceedings of the Fourteenth International Joint Conference of Artificial Intelligence (IJCAI '95)*, pages 1534–1540, 1995.
- [20] Paolo Liberatore and Marco Schaerf. Arbitration: A commutative operator for belief revision. In *Proceedings of the Second World Conference on the Fundamentals of Artificial Intelligence (WOCFAI '95)*, pages 217–228, 1995.
- [21] Abhaya C. Nayak. Iterated belief change based on epistemic entrenchment. *Erkenntnis*, 41:353–390, 1994.
- [22] Abhaya C. Nayak, Norman Y. Foo, Maurice Pagnucco, and Abdul Sattar. Changing conditional beliefs unconditionally. In *Proceedings of the Sixth Conference on Theoretical Aspects of Rationality and Knowledge (TARK VI)*, pages 119–135, 1996.
- [23] Peter Z. Revesz. On the semantics of theory change: Arbitration between old and new information. In *Proceedings of the Twelfth ACM SIGACT SIGMOD SIGART Symposium on Principles of Database Systems (PODS '93)*, pages 71–82, 1993.
- [24] Erik Sandewall and Yoav Shoham. Non-monotonic temporal reasoning. In D. Gabbay, C. J. Hogger, and J. A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 4, pages 439–498. Oxford University Press, 1995.
- [25] Wolfgang Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In William L. Harper and Brian Skyrms, editors, *Causation in Decision, Belief Change, and Statistics, II*, pages 105–134. Kluwer Academic Publishers, 1988.

- [26] Mary-Anne Williams. Transmutations of knowledge systems. In *Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning (KR '94)*, pages 619–629, 1994.

## A Appendix: Proofs

**Proposition 1**  $\preceq$  is a total pre-order on  $\mathcal{W}$ .

**Proof:** In the following, we use  $s_{ij}$  to denote the source  $\max(\Psi(w_i, w_j))$ . By Definition 2,  $\Psi(w_i, w_j)$  is always defined and non-empty, so  $s_{ij}$  is always defined.

We need to show that  $\preceq$  is connected and transitive; we first prove the former. Suppose  $w_1, w_2 \in \mathcal{W}$ .  $\sqsubseteq$  is a total order on  $\mathcal{S}$ , so either  $s_{21} \not\sqsubseteq s_{12}$  or  $s_{12} \not\sqsubseteq s_{21}$ . Thus, by Definition 3, either  $\Psi_{\sqsubseteq}(w_1, w_2) = s_0$  or  $\Psi_{\sqsubseteq}(w_2, w_1) = s_0$ , respectively. By Definition 4, either  $w_2 \preceq w_1$  or  $w_1 \preceq w_2$ .

Now we show that  $\preceq$  is transitive. First, we make the following observations for arbitrary  $w_i, w_j \in \mathcal{W}$  and  $S \subset \mathcal{S}$ :

1.  $w_i \preceq w_j$  iff  $s_{ji} \sqsubseteq s_{ij}$ , by a straightforward application of Definitions 4 and 3 show this.
2.  $w_i \leq_{s_{ij}} w_j$ , since by Definition 2 either  $s_{ij} = s_0$  and, therefore, fully connected, or  $w_i <_{s_{ij}} w_j$ , in which case the result must be true given  $s_{ij}$  is connected.
3. If  $s_{ij} = s_0$  then  $\forall s \in S. w_j \leq_s w_i$ . If there was an  $s'$  such that this were false, then  $s'$  would be in  $\Psi(w_i, w_j)$ , and since  $s_0$  is minimal wrt  $\sqsubseteq$ ,  $s_{ij}$  would not be  $\max(\Psi(w_i, w_j))$ , a contradiction.
4. If  $s_{ij} = s_{ji}$  then  $s_{ij} = s_0$ . If not, then  $w_i <_{s_{ij}} w_j$  by Definition 2 and, since  $s_{ij} = s_{ji}$ ,  $w_j <_{s_{ij}} w_i$ , a contradiction since  $\leq_{s_{ij}}$  is connected.
5. If  $s_{ij} \sqsubset s_{kl}$ , then  $w_j \leq_{s_{kl}} w_i$ . If not, then  $s_{kl} \in \Psi(w_i, w_j)$  and  $s_{ij} \neq \max(\Psi(w_i, w_j))$ , a contradiction.

Now, suppose  $w_1, w_2, w_3 \in \mathcal{W}$ ,  $w_1 \preceq w_2$ , and  $w_2 \preceq w_3$ . We need to show that  $w_1 \preceq w_3$ . By the first Observation 1,  $s_{21} \sqsubseteq s_{12}$  and  $s_{32} \sqsubseteq s_{23}$ , and it suffices to establish that  $s_{31} \sqsubseteq s_{13}$ . If  $s_{31} = s_0$  then we're done. Assume not. Then  $w_3 <_{s_{31}} w_1$  by Definition 2.

**Case 1:**  $s_{21} = s_{12}$ . Then  $s_{21} = s_0$  (Observation 4) which implies  $\forall s \in S. w_1 \leq_s w_2$ . In particular,  $w_1 \leq_{s_{31}} w_2$ , so  $w_3 <_{s_{31}} w_2$  which implies  $s_{13} \in \Psi(w_3, w_2)$  by Definition 2. Thus,  $s_{31} \sqsubseteq s_{32}$ , so  $s_{31} \sqsubseteq s_{23}$ . Since  $s_{31} \neq s_0$  and  $s_0$  is minimal wrt  $\sqsubseteq$ ,  $s_{23} \neq s_0$ . Consequently,  $w_2 <_{s_{23}} w_3$ , and since  $w_1 \leq_{s_{23}} w_2$ ,  $w_1 <_{s_{23}} w_3$ , so  $s_{23} \in \Psi(w_1, w_3)$ . Therefore,  $s_{23} \sqsubseteq s_{13}$ , and by transitivity,  $s_{31} \sqsubseteq s_{13}$ .

**Case 2:**  $s_{32} = s_{23}$ . The proof that  $s_{31} \sqsubseteq s_{13}$  is almost identical to the first case, switching  $s_{32}$  with  $s_{21}$  and  $s_{23}$  with  $s_{12}$ .

**Case 3:**  $s_{21} \sqsubset s_{12}$  and  $s_{32} \sqsubset s_{23}$ . We prove the result by contradiction. Suppose  $s_{13} \sqsubset s_{31}$ . Then  $s_{31} \neq s_0$  and  $w_3 <_{s_{31}} w_1$  by Definition 2.

First suppose  $s_{21} = s_0$ . Then, by Observation 3,  $\forall s \in S. w_1 \leq_s w_2$  and, in particular,  $w_1 \leq_{s_{31}} w_2$  and  $w_1 \leq_{s_{23}} w_2$ . Note that  $s_{23} \neq s_0$ , so  $w_2 <_{s_{23}} w_3$  by

Definition 2. Thus,  $w_3 <_{s_{31}} w_2$  and  $w_1 <_{s_{23}} w_3$ , implying that  $s_{31} \in \Psi(w_3, w_2)$  and  $s_{23} \in \Psi(w_1, w_3)$ . So, by  $s_{31} \sqsubseteq s_{32} \sqsubseteq s_{23} \sqsubseteq s_{13}$ , a contradiction. Suppose instead  $s_{21} \neq s_0$ . Then  $w_2 <_{s_{21}} w_1$ . We consider two cases:

**Case 3a:**  $s_{12} \sqsubseteq s_{32}$ . Then  $s_{21} \sqsubset s_{23}$ , so by Observation 5,  $w_1 \leq_{s_{23}} w_2$ . Since  $w_2 \leq_{s_{23}} w_3$  by Observation 2,  $w_1 <_{s_{23}} w_3$ , so  $s_{23} \in \Psi(w_1, w_3)$ . Thus,  $s_{23} \sqsubseteq s_{13}$  and by transitivity of  $\sqsubseteq$ ,  $s_{21} \sqsubset s_{31}$ . By Observation 5,  $w_1 \leq_{s_{31}} w_2$ . So  $w_3 <_{s_{31}} w_2$  by transitivity of  $\leq_{s_{31}}$ . Therefore,  $s_{31} \in \Psi(w_3, w_2)$ , so  $s_{31} \sqsubseteq s_{32}$ . But then  $s_{31} \sqsubset s_{23} \sqsubseteq s_{13}$ , contradicting our assumption.

**Case 3b:**  $s_{32} \sqsubset s_{12}$ . Then  $s_{12} \neq s_0$  and  $w_2 \leq_{s_{12}} w_3$  by Observation 5.  $s_{12} \neq s_0$  implies  $w_1 <_{s_{12}} w_2$  by Definition 2, so  $w_1 <_{s_{12}} w_3$ . Thus,  $s_{12} \in \Psi(w_1, w_3)$ , implying that  $s_{12} \sqsubseteq s_{13}$ . Thus, by transitivity of  $\sqsubseteq$ ,  $s_{21} \sqsubset s_{31}$  and, by Observation 5,  $w_1 \leq_{s_{31}} w_2$ . By transitivity of  $\leq_{s_{31}}$ ,  $w_3 <_{s_{31}} w_2$ , so  $s_{31} \in \Psi(w_3, w_2)$  which implies  $s_{31} \sqsubseteq s_{32}$ . But then  $s_{31} \sqsubset s_{12} \sqsubseteq s_{13}$ , contradicting our assumption.

■

### Proposition 2

1.  $(\Psi_1 \odot \Psi_2)(w_1, w_2) = \Psi_1(w_1, w_2) \cup \Psi_2(w_1, w_2)$
2.  $(\Psi_1 \odot \Psi_2)_{\sqsubseteq}(w_1, w_2) =$   
 $\max(\Psi_{1_{\sqsubseteq}}(w_1, w_2), \Psi_{2_{\sqsubseteq}}(w_1, w_2))$   
*if*  $\max(\Psi_{1_{\sqsubseteq}}(w_2, w_1), \Psi_{2_{\sqsubseteq}}(w_2, w_1)) \sqsubset$   
 $\max(\Psi_{1_{\sqsubseteq}}(w_1, w_2), \Psi_{2_{\sqsubseteq}}(w_1, w_2))$ , and  
 $s_0$  otherwise.
3. If  $\preceq_1, \preceq_2$ , and  $\preceq$  are the orderings induced by  $\Psi_{1_{\sqsubseteq}}, \Psi_{2_{\sqsubseteq}}$ , and  $(\Psi_1 \odot \Psi_2)_{\sqsubseteq}$ , respectively, then  
 $w_1 \prec w_2$  iff  
 $w_1 \prec_1 w_2$  and  $\max(\Psi_{2_{\sqsubseteq}}(w_1, w_2), \Psi_{2_{\sqsubseteq}}(w_2, w_1)) \sqsubseteq \Psi_{1_{\sqsubseteq}}(w_1, w_2)$  or  
 $w_1 \prec_2 w_2$  and  $\max(\Psi_{1_{\sqsubseteq}}(w_1, w_2), \Psi_{1_{\sqsubseteq}}(w_2, w_1)) \sqsubseteq \Psi_{2_{\sqsubseteq}}(w_1, w_2)$ .

### Proof:

1. Suppose  $w_1, w_2 \in \mathcal{W}$ , and  $\Psi_1$  and  $\Psi_2$  are induced by sets of sources  $S_1, S_2 \in \mathcal{S}$ , respectively. Suppose  $s \in (\Psi_1 \odot \Psi_2)(w_1, w_2)$ . Then, by Definitions 2 and 5,

$$\begin{aligned} s &\in \{(\mathcal{W}, \leq) \in S_1 \cup S_2 : w_1 < w_2\} \cup \{s_0\} \\ &= \{(\mathcal{W}, \leq) \in S_1 : w_1 < w_2\} \cup \{s_0\} \cup \\ &\quad \{(\mathcal{W}, \leq) \in S_2 : w_1 < w_2\} \cup \{s_0\} \\ &= \Psi_1(w_1, w_2) \cup \Psi_2(w_1, w_2) \end{aligned}$$

Now suppose  $s \in \Psi_1(w_1, w_2) \cup \Psi_2(w_1, w_2)$ . Then, again applying Definitions 2 and 5,

$$\begin{aligned} s &\in (\{(\mathcal{W}, \leq) \in S_1 : w_1 < w_2\} \cup \{s_0\}) \cup \\ &\quad (\{(\mathcal{W}, \leq) \in S_2 : w_1 < w_2\} \cup \{s_0\}) \\ &= \{(\mathcal{W}, \leq) \in S_1 \cup S_2 : w_1 < w_2\} \cup \{s_0\} \\ &= (\Psi_1 \odot \Psi_2)(w_1, w_2) \end{aligned}$$

2. By Definition 3,

$$\begin{aligned} (\Psi_1 \odot \Psi_2)_{\sqsubset}(w_1, w_2) &= \\ \max((\Psi_1 \odot \Psi_2)(w_1, w_2)) & \\ \text{if } \max((\Psi_1 \odot \Psi_2)(w_2, w_1)) \sqsubset & \\ \max((\Psi_1 \odot \Psi_2)(w_1, w_2)), \text{ and} & \\ s_0 \text{ otherwise.} & \end{aligned}$$

Thus, it suffices to show that

$$\begin{aligned} \max(\Psi_{1\sqsubset}(w_2, w_1), \Psi_{2\sqsubset}(w_2, w_1)) & \\ \sqsubset \max(\Psi_{1\sqsubset}(w_2, w_1), \Psi_{2\sqsubset}(w_2, w_1)) & \\ \text{iff} & \\ \max((\Psi_1 \odot \Psi_2)(w_2, w_1)) & \\ \sqsubset \max((\Psi_1 \odot \Psi_2)(w_1, w_2)) & \\ = \max(\Psi_{1\sqsubset}(w_1, w_2), \Psi_{2\sqsubset}(w_1, w_2)) & \end{aligned}$$

Now, by the first part of this proposition,

$$(\Psi_1 \odot \Psi_2)(w_1, w_2) = \Psi_1(w_1, w_2) \cup \Psi_2(w_1, w_2)$$

so

$$\begin{aligned} \max((\Psi_1 \odot \Psi_2)(w_1, w_2)) & \\ = \max(\Psi_1(w_1, w_2) \cup \Psi_2(w_1, w_2)) & \\ = \max(\max(\Psi_1(w_1, w_2)), \max(\Psi_2(w_1, w_2))) & \end{aligned}$$

and, similarly,

$$\begin{aligned} \max((\Psi_1 \odot \Psi_2)(w_2, w_1)) & \\ = \max(\max(\Psi_1(w_2, w_1)), \max(\Psi_2(w_2, w_1))) & \end{aligned}$$

( $\Leftarrow$ ) Suppose

$$\begin{aligned} \max((\Psi_1 \odot \Psi_2)(w_2, w_1)) & \\ \sqsubset \max((\Psi_1 \odot \Psi_2)(w_1, w_2)) & \\ = \max(\Psi_{1\sqsubset}(w_1, w_2), \Psi_{2\sqsubset}(w_1, w_2)) & \end{aligned}$$

Then

$$\begin{aligned} & \max((\Psi_1 \odot \Psi_2)(w_2, w_1)) \\ &= \max(\max(\Psi_1(w_2, w_1)), \max(\Psi_2(w_2, w_1))) \\ &\sqsubseteq \max(\Psi_{1\square}(w_1, w_2), \Psi_{2\square}(w_1, w_2)) \end{aligned}$$

so it is enough to show

$$\begin{aligned} & \max(\Psi_{1\square}(w_2, w_1), \Psi_{2\square}(w_2, w_1)) \\ &\sqsubseteq \max(\max(\Psi_1(w_2, w_1)), \max(\Psi_2(w_2, w_1))) \end{aligned}$$

To do so, we only need to show that  $\Psi_{1\square}(w_2, w_1) \sqsubseteq \max(\Psi_1(w_2, w_1))$  and  $\Psi_{2\square}(w_2, w_1) \sqsubseteq \max(\Psi_2(w_2, w_1))$ . These follow immediately from Definition 3 and the fact that  $s_0$  is minimal wrt  $\sqsubseteq$ .

( $\implies$ ) Suppose

$$\begin{aligned} & \max(\Psi_{1\square}(w_2, w_1), \Psi_{2\square}(w_2, w_1)) \\ &\sqsubseteq \max(\Psi_{1\square}(w_1, w_2), \Psi_{2\square}(w_1, w_2)). \end{aligned}$$

Assume, without loss of generality, that  $\Psi_{2\square}(w_1, w_2) \sqsubseteq \Psi_{1\square}(w_1, w_2)$ . Then,  $\Psi_{1\square}(w_2, w_1) \sqsubset \Psi_{1\square}(w_1, w_2)$  and since  $\Psi_{1\square}(w_1, w_2) \neq s_0$ , by Definition 3  $\max(\Psi_1(w_1, w_2)) = \Psi_{1\square}(w_1, w_2)$ . Observe also that  $\Psi_{2\square}(w_2, w_1) \sqsubset \Psi_{1\square}(w_1, w_2)$ .

We now show that  $\max(\Psi_2(w_1, w_2)) \sqsubseteq \Psi_{1\square}(w_1, w_2)$ . Suppose  $\max(\Psi_2(w_2, w_1)) \sqsubset \max(\Psi_2(w_1, w_2))$ . Then by Definition 3,

$$\begin{aligned} & \max(\Psi_2(w_1, w_2)) \\ &= \Psi_{2\square}(w_1, w_2) \\ &\sqsubseteq \Psi_{1\square}(w_1, w_2) \end{aligned}$$

On the other hand, if  $\max(\Psi_2(w_2, w_1)) = \max(\Psi_2(w_1, w_2))$  then  $\max(\Psi_2(w_1, w_2)) = s_0 \sqsubset \Psi_{1\square}(w_1, w_2)$ . Finally, if  $\max(\Psi_2(w_1, w_2)) \sqsubset \max(\Psi_2(w_2, w_1))$  then from Definition 3 and the observation above,

$$\begin{aligned} & \max(\Psi_2(w_1, w_2)) \\ &\sqsubset \max(\Psi_2(w_2, w_1)) \\ &= \Psi_{2\square}(w_2, w_1) \\ &\sqsubset \Psi_{1\square}(w_1, w_2) \end{aligned}$$

Therefore,

$$\begin{aligned} & \max((\Psi_1 \odot \Psi_2)(w_1, w_2)) \\ &= \max(\max(\Psi_1(w_1, w_2)), \max(\Psi_2(w_1, w_2))) \\ &= \max(\Psi_{1\square}(w_1, w_2), \max(\Psi_2(w_1, w_2))) \\ &= \Psi_{1\square}(w_1, w_2). \end{aligned}$$

Note that we have also shown that  $\max(\Psi_2(w_2, w_1)) \sqsubset \Psi_{1\sqsubset}(w_1, w_2)$ .

Now, since  $\Psi_{1\sqsubset}(w_2, w_1) \sqsubset \Psi_{1\sqsubset}(w_1, w_2)$  and, consequently,  $\Psi_{1\sqsubset}(w_1, w_2) \neq s_0$ , by Definition 3  $\max(\Psi_1(w_2, w_1)) \sqsubset \max(\Psi_1(w_1, w_2)) = \Psi_{1\sqsubset}(w_1, w_2)$ . Putting this together with the results from the previous paragraph, we have that

$$\begin{aligned} & \max((\Psi_1 \otimes \Psi_2)(w_2, w_1)) \\ &= \max(\max(\Psi_1(w_2, w_1)), \max(\Psi_2(w_2, w_1))) \\ &\sqsubset \Psi_{1\sqsubset}(w_1, w_2) \\ &= \max((\Psi_1 \otimes \Psi_2)(w_1, w_2)) \end{aligned}$$

and, given our earlier assumption,

$$\begin{aligned} & \max((\Psi_1 \otimes \Psi_2)(w_1, w_2)) \\ &= \max(\Psi_{1\sqsubset}(w_1, w_2), \Psi_{2\sqsubset}(w_1, w_2)). \end{aligned}$$

■

3. We start by proving an auxiliary lemma.

**Lemma 1** *Given  $\mathcal{W}$  and  $\mathcal{S}$  as above, for any two pedigreed belief states  $\Psi_1$ ,  $\Psi_2$ , and any two worlds  $w_1, w_2 \in \mathcal{W}$ ,*

$$\Psi_1(w_1, w_2) \cap \Psi_2(w_2, w_1) = \{s_0\}.$$

**Proof:** By Definition 2,  $s_0 \in \Psi_1(w_1, w_2)$  and  $s_0 \in \Psi_2(w_2, w_1)$ , so  $\{s_0\} \subseteq \Psi_1(w_1, w_2) \cap \Psi_2(w_2, w_1)$ .

Suppose  $s = (\mathcal{W}, \leq_s) \in \Psi_1(w_1, w_2) \cap \Psi_2(w_2, w_1)$  for some  $s \in \mathcal{S}$ . Then  $s \in \Psi_1(w_1, w_2)$  and  $s \in \Psi_2(w_2, w_1)$ . Assume  $s \neq s_0$ . Then, by Definition 2,  $w_1 <_s w_2$  and  $w_1 <_s w_2$ , a contradiction. Therefore,  $s = s_0$ , so  $\Psi_1(w_1, w_2) \cap \Psi_2(w_2, w_1) \subseteq \{s_0\}$ . ■

**Corollary 2.1** *Given  $\mathcal{W}$ ,  $\mathcal{S}$ ,  $\Psi_1$ ,  $\Psi_2$ ,  $w_1$ , and  $w_2$  as above,  $\max(\Psi_1(w_1, w_2)) = \max(\Psi_2(w_2, w_1))$  implies  $\max(\Psi_1(w_1, w_2)) = s_0$ .*

**Proof:** Suppose  $\max(\Psi_1(w_1, w_2)) = \max(\Psi_2(w_2, w_1))$ . Then, since  $\max(\Psi_1(w_1, w_2)) \in \Psi_1(w_1, w_2)$  and  $\max(\Psi_2(w_2, w_1)) \in \Psi_2(w_2, w_1)$ ,  $\max(\Psi_1(w_1, w_2)) \in \Psi_1(w_1, w_2) \cap \Psi_2(w_2, w_1)$ . By the above result,  $\max(\Psi_1(w_1, w_2)) = s_0$ . ■

We proceed to prove the proposition.

( $\Leftarrow$ ) Suppose  $w_1 \prec_1 w_2$  and  $\max(\Psi_{2\sqsubset}(w_1, w_2), \Psi_{2\sqsubset}(w_2, w_1)) \sqsubseteq \Psi_{1\sqsubset}(w_1, w_2)$ . Then, by Definition 4,  $\Psi_{1\sqsubset}(w_1, w_2) \neq s_0$  and  $\Psi_{1\sqsubset}(w_2, w_1) = s_0$ , so

$\Psi_{1\square}(w_2, w_1) \sqsubset \Psi_{1\square}(w_1, w_2)$ . Also,  $\Psi_{2\square}(w_1, w_2) \sqsubseteq \Psi_{1\square}(w_1, w_2)$  and  $\Psi_{2\square}(w_2, w_1) \sqsubseteq \Psi_{1\square}(w_1, w_2)$ . Suppose  $\Psi_{2\square}(w_2, w_1) = \Psi_{1\square}(w_1, w_2)$ . Then,  $\Psi_{2\square}(w_2, w_1) \neq s_0$  and, thus,  $\Psi_{2\square}(w_1, w_2) = s_0$ . Furthermore, by Definition 3,  $\Psi_{2\square}(w_2, w_1) = \max(\Psi_2(w_2, w_1))$  so

$$\begin{aligned} & \max(\Psi_1(w_1, w_2)) \\ &= \Psi_{1\square}(w_1, w_2) \\ &= \Psi_{2\square}(w_2, w_1) \\ &= \max(\Psi_2(w_2, w_1)) \end{aligned}$$

and, by Corollary 2.1,  $\max(\Psi_1(w_1, w_2)) = s_0$ , a contradiction. Consequently,  $\Psi_{2\square}(w_2, w_1) \sqsubset \Psi_{1\square}(w_1, w_2)$ , so

$$\begin{aligned} & \max(\Psi_{1\square}(w_2, w_1), \Psi_{2\square}(w_2, w_1)) \\ & \sqsubset \Psi_{1\square}(w_1, w_2) \\ &= \max(\Psi_{1\square}(w_1, w_2), \Psi_{2\square}(w_1, w_2)). \end{aligned}$$

By Definition 3 and Proposition 2,

$$\begin{aligned} (\Psi_1 \oslash \Psi_2)_{\square}(w_1, w_2) &\neq s_0 \text{ and} \\ (\Psi_1 \oslash \Psi_2)_{\square}(w_2, w_1) &= s_0 \end{aligned}$$

so, by Definition 4,  $w_1 \prec w_2$ .

Similarly, if  $w_1 \prec_2 w_2$  and  $\max(\Psi_{1\square}(w_1, w_2), \Psi_{1\square}(w_2, w_1)) \sqsubseteq \Psi_{2\square}(w_1, w_2)$ , then  $w_1 \prec w_2$ .

( $\Rightarrow$ ) Suppose  $w_1 \prec w_2$ . By Definition 4,

$$\begin{aligned} (\Psi_1 \oslash \Psi_2)_{\square}(w_1, w_2) &\neq s_0 \text{ and} \\ (\Psi_1 \oslash \Psi_2)_{\square}(w_2, w_1) &= s_0, \end{aligned}$$

by Definition 3

$$\begin{aligned} & \max((\Psi_1 \oslash \Psi_2)(w_2, w_1)) \\ & \sqsubset \max((\Psi_1 \oslash \Psi_2)(w_1, w_2)), \end{aligned}$$

and by Proposition 2

$$\begin{aligned} & \max(\max(\Psi_1(w_2, w_1)), \max(\Psi_2(w_2, w_1))) \\ &= \max(\max(\Psi_1(w_2, w_1)) \cup \max(\Psi_2(w_2, w_1))) \\ & \sqsubset \max(\Psi_1(w_1, w_2) \cup \Psi_2(w_1, w_2)) \\ &= \max(\max(\Psi_1(w_1, w_2)), \max(\Psi_2(w_1, w_2))) \end{aligned}$$

Assume  $\max(\Psi_2(w_1, w_2)) \sqsubseteq \max(\Psi_1(w_1, w_2))$ . Then

$$\begin{aligned}\max(\Psi_1(w_2, w_1)) &\sqsubset \max(\Psi_1(w_1, w_2)) \text{ and} \\ \max(\Psi_2(w_2, w_1)) &\sqsubset \max(\Psi_1(w_1, w_2)).\end{aligned}$$

Definition 3 gives us that

$$\begin{aligned}\Psi_{1\square}(w_1, w_2) &= \max(\Psi_1(w_1, w_2)) \text{ and} \\ \Psi_{1\square}(w_2, w_1) &= s_0\end{aligned}$$

Thus, by Definition 4  $w_1 \prec_1 w_2$ . Also, from Definition 3 we know  $\Psi_{2\square}(w_1, w_2) \sqsubseteq \max(\Psi_2(w_1, w_2))$  and  $\Psi_{2\square}(w_2, w_1) \sqsubseteq \max(\Psi_2(w_2, w_1))$ . Therefore, given that our original assumption and the fact that  $\max(\Psi_2(w_2, w_1)) \sqsubset \max(\Psi_1(w_1, w_2))$ , we have

$$\begin{aligned}\max(\Psi_{2\square}(w_1, w_2), \Psi_{2\square}(w_2, w_1)) &\\ \sqsubseteq \max(\max(\Psi_2(w_1, w_2)), \max(\Psi_2(w_2, w_1))) &\\ \sqsubseteq \max(\Psi_1(w_1, w_2)) &\\ = \Psi_{1\square}(w_1, w_2) &\end{aligned}$$

Similarly, if  $\max(\Psi_1(w_1, w_2)) \sqsubseteq \max(\Psi_2(w_1, w_2))$ , then  $w_1 \prec_2 w_2$  and

$$\max(\Psi_{1\square}(w_1, w_2), \Psi_{1\square}(w_2, w_1)) \sqsubseteq \Psi_{2\square}(w_1, w_2).$$

■

**Proposition 3** Let  $\Psi_1$  and  $\Psi_2$  be pedigreed belief states, and let  $\preceq_1$  and  $\preceq_2$  be the orderings induced by  $\Psi_{1\square}$  and  $\Psi_{2\square}$ , respectively. Further, let  $\preceq$  be the ordering induced by  $(\Psi_1 \odot \Psi_2)_{\square}$ . If  $\Psi_2 \geq \Psi_1$ , then  $w_1 \prec_2 w_2$  implies  $w_1 \prec w_2$  for all  $w_1, w_2 \in \mathcal{W}$ .

**Proof:** Suppose  $\Psi_1 \geq \Psi_2$ ,  $w_1, w_2 \in \mathcal{W}$ , and  $w_1 \preceq_2 w_2$ . By Definition 4, it suffices to show that  $(\Psi_1 \odot \Psi_2)_{\square}(w_1, w_2) \neq s_0$  and  $(\Psi_1 \odot \Psi_2)_{\square}(w_2, w_1) = s_0$ . By Definition 4,  $\Psi_{2\square}(w_1, w_2) \neq s_0$  and  $\Psi_{2\square}(w_2, w_1) = s_0$ . Since  $\Psi_1 \geq \Psi_2$  and  $\Psi_{2\square}(w_1, w_2) \neq s_0$ ,  $\max(\Psi_{1\square}(w_1, w_2), \Psi_{1\square}(w_2, w_1)) \sqsubseteq \Psi_{2\square}(w_1, w_2)$  by Definition 6. Thus,  $\Psi_{1\square}(w_1, w_2) \sqsubseteq \Psi_{2\square}(w_1, w_2)$  and  $\Psi_{1\square}(w_2, w_1) \sqsubseteq \Psi_{2\square}(w_1, w_2)$ , so  $\max(\Psi_{1\square}(w_1, w_2), \Psi_{2\square}(w_1, w_2)) = \Psi_{2\square}(w_1, w_2)$ . Also, since  $\forall s \in \mathcal{S}. s \neq s_0 \Rightarrow s_0 \sqsubset s$ , we have  $\max(\Psi_{1\square}(w_2, w_1), \Psi_{2\square}(w_2, w_1)) = \Psi_{1\square}(w_2, w_1)$ .

Suppose  $\Psi_{1\square}(w_2, w_1) = \Psi_{2\square}(w_1, w_2) = s$  for some  $s \in \mathcal{S}$ . Then, by Definition 3,  $s \in \Psi_1(w_2, w_1)$  and  $s \in \Psi_2(w_1, w_2)$ . But then, since  $s \neq s_0$ , by Definition 2  $w_2 <_s w_1$  and  $w_1 <_s w_2$ . This is a contradiction since  $s \in \mathcal{S}$  implies  $\leq_s$  is connected. Therefore,  $\max(\Psi_{1\square}(w_2, w_1), \Psi_{2\square}(w_2, w_1)) = \Psi_{1\square}(w_2, w_1) \sqsubset \Psi_{2\square}(w_1, w_2) = \max(\Psi_{1\square}(w_1, w_2), \Psi_{2\square}(w_1, w_2))$ . By Proposition 2,  $(\Psi_1 \odot \Psi_2)_{\square}(w_1, w_2) \neq s_0$  and  $(\Psi_1 \odot \Psi_2)_{\square}(w_2, w_1) = s_0$ . ■

**Proposition 4** Let  $\Psi_1$  and  $\Psi_2$  be pedigreed belief states such that  $\Psi_2 \geq \Psi_1$ . Then  $(\Psi_1 \otimes \Psi_2)\downarrow = \Psi_1 \circ (\Psi_2\downarrow)$ .

**Proof:** Let  $\preceq_1$ ,  $\preceq_2$ , and  $\preceq$  be the orderings induced by  $\Psi_1$ ,  $\Psi_2$ , and  $\Psi = \Psi_1 \otimes \Psi_2$ , respectively. Suppose  $w \in \Psi\downarrow$ . To show that  $w \in \Psi_1 \circ (\Psi_2\downarrow)$ , it suffices to first show that  $w \in \Psi_2\downarrow$  and that for every  $w' \in \Psi_2\downarrow$ ,  $w \preceq_1 w'$ . By definition,  $w \in \min(\preceq, \mathcal{W})$ , so  $\forall w' \in \mathcal{W}. w \preceq w'$ . By Proposition 3,  $\forall w' \in \mathcal{W}. w \preceq_2 w'$ . Thus,  $w \in \min(\preceq_2, \mathcal{W})$ , so  $w \in \Psi_2\downarrow$ .

Now let  $w' \in \Psi_2\downarrow$ . We show that  $w \preceq_1 w'$ . Suppose not, i.e.,  $w' \prec_1 w$ . Then, by Definition 4,  $\Psi_{1\sqsubset}(w, w') = s_0$  and  $\Psi_{1\sqsubset}(w', w) \neq s_0$ . Since  $w' \in \Psi_2\downarrow = \min(\preceq_2, \mathcal{W})$ ,  $w \preceq_2 w'$  and  $w' \preceq_2 w$ . So  $\Psi_{2\sqsubset}(w, w') = \Psi_{2\sqsubset}(w', w) = s_0$  by Definition 4. Thus,  $\max(\Psi_{2\sqsubset}(w, w'), \Psi_{2\sqsubset}(w', w)) = s_0 \sqsubset \Psi_{1\sqsubset}(w', w) = s$  since  $\forall s \in \mathcal{S}. s \neq s_0 \Rightarrow s_0 \sqsubset s$ . By Proposition 2,  $w' \prec w$ . But then  $w \notin \min(\preceq, \mathcal{W})$ . Contradiction. Therefore,  $\forall w' \in \mathcal{W}. w \preceq_1 w'$ .

We now prove the other direction of the proposition. Suppose  $w \in \Psi_1 \circ (\Psi_2\downarrow)$ . Then  $w \in \min(\preceq_1, \min(\preceq_2, \mathcal{W}))$ . This implies that  $w \in \min(\preceq_2, \mathcal{W})$  which, in turn, implies that  $\forall w' \in \mathcal{W}. w \preceq_2 w'$ . Suppose  $w' \in \mathcal{W}$ . We show that  $w \preceq w'$ . Suppose not, i.e.,  $w' \prec w$ . Proposition 2 gives us two case:

1.  $w' \prec_2 w$ . Then  $w \notin \min(\preceq_2, \mathcal{W})$ . Contradiction.
2.  $w' \prec_1 w$ . Since  $w' \prec w$ ,  $w' \preceq_2 w$  by Proposition 3. Thus, since  $w \in \min(\preceq_2, \mathcal{W})$ , so is  $w'$ . But if  $w' \prec_1 w$ , then  $w \notin \min(\preceq_1, \min(\preceq_2, \mathcal{W}))$ . Contradiction.

Therefore,  $\forall w' \in \mathcal{W}. w \preceq w'$ , so  $w \in (\Psi_1 \otimes \Psi_2)\downarrow$ . ■

**Corollary 4.1** Let  $\Psi_1, \Psi_2, \Psi_3$  be pedigreed belief states such that  $\Psi_2 \geq \Psi_1$ ,  $\Psi_3 \geq \Psi_1$ , and  $\Psi_2\downarrow = \Psi_3\downarrow$ . Then  $(\Psi_1 \otimes \Psi_2)\downarrow = (\Psi_1 \otimes \Psi_3)\downarrow$ .

**Proof:** Appealing to Proposition 4,

$$\begin{aligned} (\Psi_1 \otimes \Psi_2)\downarrow &= \Psi_1 \circ (\Psi_2\downarrow) \\ &= \Psi_1 \circ (\Psi_3\downarrow) \\ &= (\Psi_1 \otimes \Psi_3)\downarrow. \end{aligned}$$

We introduce some notation for the proofs that follow: Given a belief state  $(\mathcal{W}, \leq)$ , let  $\leftarrow$  be a total order over subsets of  $\mathcal{W}$  such that if  $W, W' \subseteq \mathcal{W}$ ,  $W \leftarrow W'$  iff

1.  $W$  and  $W'$  are non-empty,

2.  $W$  and  $W'$  are (not necessarily maximal) equivalence sets, i.e., for every  $w \in W$  and  $w' \in W$ , if  $w' \in W$  then  $w \leq w'$  and  $w' \leq w$  (and similarly for  $W'$ ), and
3. worlds in  $W$  are strictly preferred to worlds in  $W'$ , i.e., for every  $w_1 \in W$  and  $w_2 \in W'$ ,  $w_1 < w_2$ .

Thus, we can represent the belief states in Figure 3 as  $A = \|p\| \leftarrow \|\neg p \wedge r\| \leftarrow \|\neg p \wedge \neg r\|$ ,  $B = \|r\| \leftarrow \|\neg p \wedge \neg r\| \leftarrow \|p \wedge \neg r\|$ , and  $C = \|r\| \leftarrow \|\neg r\|$ . AGM requires that  $(A \diamond (B \diamond C \downarrow)) \downarrow = \|\neg p \wedge \neg r\|$ . In each of the following proofs, we show that left association gives an inconsistent result, specifically,  $\|p \wedge \neg r\|$ .

**Boutilier's natural revision** The natural revision operator  $\diamond_B$  is defined as follows:

**Definition 8** If  $M = (\mathcal{W}, \leq)$  is a belief state, then  $(M \diamond_B p) = (\mathcal{W}, \leq')$  is the belief state resulting from the natural revision of  $M$  by sentence  $p$  if and only if for all  $w_1, w_2 \notin \min(\leq, p)$ ,  $w_1 \leq' w_2$  iff  $w_1 \leq w_2$  and, by the AGM postulates, for all  $w_1 \in \min(\leq, p)$  and  $w_2 \in \mathcal{W}$ ,  $w_1 \leq' w_2$ .

**Proposition 5** The resulting belief sets using left and right association of Boutilier's natural revision operators can be inconsistent.

**Proof:** Applying the operator to the belief states  $A, B, C$ , we get  $((A \diamond_B B \downarrow) \diamond_B C \downarrow) \downarrow = \|p \wedge \neg r\|$  which is inconsistent with the result of right association. ■

**Darwiche and Pearl's formulation** Let  $M = (\mathcal{W}, \leq)$  be a belief state,  $p$  be a sentence in  $\mathcal{L}$ . Darwiche and Pearl suggest a set of postulates (see [10] for their enumeration) to supplement the AGM postulates for iterated revision, then show by way of a representation theorem that an AGM operator  $\diamond_{DP}$  satisfying the postulates obeys the following rules:

1. If  $w_1 \models p$  and  $w_2 \models p$ , then  $w_1 \leq w_2$  iff  $w_1 \leq' w_2$ .
2. If  $w_1 \models \neg p$  and  $w_2 \models \neg p$ , then  $w_1 \leq w_2$  iff  $w_1 \leq' w_2$ .
3. If  $w_1 \models p$  and  $w_2 \models \neg p$ , then  $w_1 < w_2$  only if and  $w_1 <' w_2$ .
4. If  $w_1 \models p$  and  $w_2 \models \neg p$ , then  $w_1 \leq w_2$  only if  $w_1 \leq' w_2$ .

where  $(M \diamond_{DP} p) = (\mathcal{W}, \leq')$  is the result of revising  $M$  by  $p$ .

**Proposition 6** The resulting belief sets using left and right association of any revision operators satisfying the Darwiche and Pearl postulates can be inconsistent.

**Proof:** Let  $\diamond_{DP}$  be an AGM operator that is a member of the above family of operators. Then, given  $A, B, C$  as above, by the third rule  $\|p \wedge \neg r\| \leftarrow \|\neg p \wedge \neg r\|$  in  $A \diamond_B B \downarrow$ , so  $((A \diamond_B B \downarrow) \diamond_B C \downarrow) \downarrow = \|p \wedge \neg r\|$  which is inconsistent with the result of right association. ■

**Spohn's conditionalization** Let  $\mathcal{N}$  be the set of ordinals.

**Definition 9** An ordinal conditional function (OCF) is any function  $\kappa : \mathcal{W} \mapsto \mathcal{N}$  such that  $\exists w \in \mathcal{W}. \kappa(w) = 0$ . For  $W \subseteq \mathcal{W}$ , we define  $\kappa(W) = \min_{w \in W} \kappa(w)$ . The belief set of an OCF is  $\kappa\downarrow = \{w \in \mathcal{W} \mid \kappa(w) = 0\}$ .

**Definition 10** Let  $\kappa$  be an OCF.  $\diamond_\alpha$  is an  $\alpha$ -conditionalization operator iff  $\alpha$  is a non-zero ordinal and, for any sentence  $p \in \mathcal{L}$  and any  $w \in \mathcal{W}$ ,

$$(\kappa \diamond_\alpha p)(w) = \begin{cases} \kappa(w) - \kappa(p) & \text{if } w \models p \\ \alpha + \kappa(w) - \kappa(\neg p) & \text{if } w \models \neg p \end{cases}$$

**Proposition 7** The resulting belief sets using left and right association of any combination of  $\alpha$ -conditionalization operators can be inconsistent.

**Proof:** Let  $\kappa_A, \kappa_B, \kappa_C$  be the OCFs representing belief states  $A, B, C$ , respectively, such that  $\kappa_A(p \wedge r) = \kappa_A(p \wedge \neg r) = 0 < \kappa_A(\neg p \wedge r) < \kappa_A(\neg p \wedge \neg r)$ ,  $\kappa_B(p \wedge r) = \kappa_B(\neg p \wedge r) = 0 < \kappa_B(\neg p \wedge \neg r) < \kappa_B(p \wedge \neg r)$ , and  $\kappa_C(p \wedge \neg r) = \kappa_C(\neg p \wedge \neg r) = 0 < \kappa_C(p \wedge r) = \kappa_C(\neg p \wedge r)$ . Let  $\diamond_{\alpha_1}$  and  $\diamond_{\alpha_2}$  be any two  $\alpha$ -conditionalization operators. It is easily seen that  $(\kappa_A \diamond_{\alpha_1} (\kappa_B \diamond_{\alpha_2} \kappa_C\downarrow)\downarrow)\downarrow = \|\neg p \wedge \neg r\|$ . Now, by Definition 10,  $(\kappa_A \diamond_{\alpha_1} \kappa_B\downarrow)(p \wedge \neg r) < (\kappa_A \diamond_{\alpha_1} \kappa_B\downarrow)(\neg p \wedge \neg r)$ . Subsequent conditioning by  $\kappa_C\downarrow$  using the  $\alpha_2$  operator preserves this ordering and produces the belief set  $\|p \wedge \neg r\|$ . ■

**Lehmann's formulation** We refer the reader to [19] for the postulates Lehmann proposes should govern the behavior of a sequence of revisions. Lehmann gives model-theoretic semantics in terms of *widening rank models*, defined below. Using these models, he describes a recursive definition for computing the belief set that results from a sequence of revisions that obey the postulates.

**Definition 11** A widening rank model is a function  $WR : \mathcal{N} \mapsto 2^{\mathcal{W}} \setminus \emptyset$  such that

1. for any  $n, m \in \mathcal{N}$ , if  $n \leq m$  then  $WR(n) \subseteq WR(m)$ , and
2. for any  $w \in \mathcal{W}$ , there is some  $n \in \mathcal{N}$  such that  $w \in WR(n)$ ,

where  $\mathcal{N}$  is a sufficiently long initial segment of the ordinals. For  $p \in \mathcal{L}$ , we define  $rank(p) = \operatorname{argmin}_{n \in \mathcal{N}} (w \in WR(n) \wedge w \models p)$ . The belief set  $WR\downarrow = WR(0)$ .

Let  $\sigma$  be a sequence of sentences in  $\mathcal{L}$  where  $\emptyset$  is the empty sequence and  $\cdot$  is the concatenation operator.

**Definition 12** Given a widening rank model  $WR$ , the belief set resulting from the revision sequence corresponding to  $\sigma$  and obeying Lehmann's postulates, denoted  $[\sigma]_{WR}$ , is  $\pi(\sigma)$  defined recursively as follows:

1.  $r(\emptyset) = 0$  and  $\pi(\emptyset) = \mathcal{W}\mathcal{R}(0)$ .

2. If  $\tau$  is a sentence sequence,  $p \in \mathcal{L}$ , and there exists  $w \in \pi(\tau)$  such that  $w \models p$ , then  $r(\tau \cdot p) = r(\tau)$  and  $\pi(\tau \cdot p) = \{w \in \pi(\tau) \mid w \models p\}$ .
3. Otherwise,  $r(\tau \cdot p)$  is the smallest  $n > r(\tau)$  such that there exists  $w \in WR(n)$  and  $w \models p$ , and  $\pi(\tau \cdot p) = \{w \in WR(n) \mid w \models p\}$ .

where  $r$  maps revision sequences to ordinals, and  $\pi$  maps revision sequences to subsets of  $\mathcal{W}$ .

This procedure is equivalent to iteratively applying the following revision operator to the members of  $\sigma$ :

**Definition 13** If  $WR$  is a widening rank model over  $\mathcal{W}$ , then the widening rank model  $(WR \diamond_L p)$  resulting from the revision of  $WR$  by sentence  $p$  is defined as follows:

1.  $(WR \diamond_L p)(0) = \{w \in \mathcal{W} \mid w \models p \text{ and } w \in WR(rank(p))\}$ .
2. For all  $n \in \mathbb{N}$  such that  $n > 0$ ,  $(WR \diamond_L p)(n) = WR(rank(p) + n)$ .

Let  $WR_\sigma$  be the result of using  $\diamond_L$  to iteratively revise  $WR$  by consecutive members of  $\sigma$ , that is,  $WR_\emptyset = WR$  and, recursively,  $WR_{\sigma \cdot p} = WR_\sigma \diamond_L p$  for any  $p \in \mathcal{L}$ . Let  $rank_\sigma(p)$  be the rank of  $p$  in  $WR_\sigma$ .

**Lemma 2** Let  $WR$  be a widening rank model. Then  $\pi(\sigma) = WR_\sigma(0)$ , and for all  $n \in \mathbb{N}$  such that  $n > 0$ ,  $WR(r(\sigma) + n) = WR_\sigma(n)$ . In particular,  $[\sigma]_{WR} = WR_\sigma \downarrow$ .

**Proof:** The proof is by induction on the length of  $\sigma$ .

**Base case:** If  $\sigma = \emptyset$ , then  $WR = WR_\sigma$ , so

$$\begin{aligned}\pi(\sigma) \\ &= WR(0) \\ &= WR_\sigma(0).\end{aligned}$$

Furthermore, for all  $n > 0$ ,

$$\begin{aligned}WR(r(\sigma) + n) \\ &= WR(n) \\ &= WR_\sigma(n).\end{aligned}$$

**Inductive case:** Suppose  $\pi(\sigma) = WR_\sigma(0)$  and for all  $n > 0$ ,  $WR(r(\sigma) + n) = WR_\sigma(n)$ . We show that  $\pi(\sigma \cdot p) = WR_{\sigma \cdot p}(0)$  and for all  $n > 0$ ,  $WR(r(\sigma \cdot p) + n) = WR_{\sigma \cdot p}(n)$  where  $p \in \mathcal{L}$ .

First note that  $rank_\sigma(p) = r(\sigma \cdot p) - r(\sigma)$ . If  $rank_\sigma(p) = 0$ , then there exists  $w \in WR_\sigma(0)$  such that  $w \models p$ . By the inductive hypothesis,  $w \in \pi(\sigma)$  so, by Definition 12,  $r(\sigma \cdot p) = r(\sigma)$  and  $rank_\sigma(p) = r(\sigma \cdot p) - r(\sigma)$ . On the other hand,

if  $\text{rank}_\sigma(p) = n > 0$ , then there exists  $w \in \text{WR}_\sigma(n) = \text{WR}(r(\sigma) + n)$  such that  $w \models p$ . Moreover, for all  $w \in \text{WR}_\sigma(0) = \pi(\sigma)$ , and for all  $w \in \text{WR}_\sigma(n') = \text{WR}(r(\sigma) + n')$  such that  $0 < n' < n$ ,  $w \models \neg p$ . Thus,  $r(\sigma \cdot p) = r(\sigma) + n$ , so  $\text{rank}_\sigma(p) = n = r(\sigma \cdot p) - r(\sigma)$ .

Suppose  $w \in \pi(\sigma \cdot p)$ . Then  $w \models p$ . If there exists  $w' \in \pi(\sigma)$  such that  $w' \models p$ , then by Definition 12,  $r(\sigma \cdot p) = r(\sigma)$ , so  $\text{rank}_\sigma(p) = 0$  and, by Definition 13,  $w \in \text{WR}_{\sigma \cdot p}(0)$ . Otherwise,  $r(\sigma \cdot p) = n > r(\sigma)$  is the smallest ordinal such that there exists  $w' \in \text{WR}(n)$  and  $w' \models p$ , and  $\pi(\sigma \cdot p) = \{w \in \text{WR}(n) \mid w \models p\}$ . Therefore,

$$\begin{aligned} w &\in \text{WR}(r(\sigma \cdot p)) \\ &= \text{WR}(r(\sigma) + \text{rank}_\sigma(p)) \\ &= \text{WR}_\sigma(\text{rank}_\sigma(p)). \end{aligned}$$

and, by Definition 13,  $w \in \text{WR}_{\sigma \cdot p}(0)$ .

Suppose  $w \in \text{WR}_{\sigma \cdot p}(0)$ . Then  $w \models p$  and  $w \in \text{WR}_\sigma(\text{rank}_\sigma(p))$ . If  $\text{rank}_\sigma(p) = 0$ , then  $r(\sigma \cdot p) = r(\sigma)$  and

$$\begin{aligned} w &\in \text{WR}_\sigma(0) \\ &= \pi(\sigma) \\ &= \pi(\sigma \cdot p). \end{aligned}$$

Otherwise,  $\text{rank}_\sigma(p) > 0$ , so  $r(\sigma \cdot p) > r(\sigma)$  and

$$\begin{aligned} w &\in \text{WR}_\sigma(\text{rank}_\sigma(p)) \\ &= \text{WR}(r(\sigma) + \text{rank}_\sigma(p)) \\ &= \text{WR}(r(\sigma \cdot p)) \end{aligned}$$

so  $w \in \pi(\sigma \cdot p)$ . Therefore,  $\pi(\sigma \cdot p) = \text{WR}_{\sigma \cdot p}(0)$ .

Now let  $n > 0$  for some  $n \in \mathcal{N}$ . Then, since  $\text{rank}_\sigma(p) = r(\sigma \cdot p) - r(\sigma)$ , by Definition 13

$$\begin{aligned} \text{WR}(r(\sigma \cdot p) + n) \\ &= \text{WR}(r(\sigma) + \text{rank}_\sigma(p) + n) \\ &= \text{WR}_\sigma(\text{rank}_\sigma(p) + n) \\ &= \text{WR}_{\sigma \cdot p}(n) \end{aligned}$$

Finally, it follows that

$$\begin{aligned} [\sigma]_{\text{WR}} \\ &= \pi(\sigma) \\ &= \text{WR}_\sigma(0) \\ &= \text{WR}_{\sigma \downarrow}. \end{aligned}$$

■

**Proposition 8** *The resulting belief sets using left and right association of Lehmann's revision operator can be inconsistent.*

**Proof:** Widening rank models corresponding to the belief states in Figure 3 are:

$$WR_A(n) = \begin{cases} \|p\| & \text{if } n = 0 \\ \|p \vee r\| & \text{if } n = 1 \\ \mathcal{W} & \text{otherwise,} \end{cases}$$

$$WR_B(n) = \begin{cases} \|r\| & \text{if } n = 0 \\ \|\neg p \vee r\| & \text{if } n = 1 \\ \mathcal{W} & \text{otherwise,} \end{cases}$$

and

$$WR_C(n) = \begin{cases} \|\neg r\| & \text{if } n = 0 \\ \mathcal{W} & \text{otherwise} \end{cases}$$

where  $n \in \mathcal{N}$ . The reader will easily confirm that  $(WR_A \diamond_L (WR_B \diamond_L WR_C \downarrow)) \downarrow = \|\neg p \wedge \neg r\|$  whereas  $((WR_A \diamond_L WR_B \downarrow) \diamond_L WR_C \downarrow) \downarrow = \|p \wedge \neg r\|$ . ■

**Williams' transmutations** It is easy to verify that the following definition of adjustment operators is equivalent to William's.<sup>8</sup>

**Definition 14** Let  $\kappa$  be an OCF.  $\diamond_\beta$  is an  $\beta$ -adjustment operator iff  $\beta$  is a non-zero ordinal and, for any sentence  $p \in \mathcal{L}$  and any  $w \in \mathcal{W}$ ,

$$(\kappa \diamond_\beta p)(w) = \begin{cases} 0 & \text{if } w \models p \text{ and } \kappa(w) = \kappa(p) \\ \beta & \text{if } w \models \neg p, \text{ and} \\ & \kappa(w) < \beta \text{ or } \kappa(w) = \kappa(\neg p) \\ \kappa(w) & \text{otherwise.} \end{cases}$$

**Proposition 9** *The resulting belief sets using left and right association of any combination of  $\beta$ -adjustment operators can be inconsistent.*

**Proof:** Let  $\diamond_{\beta_1}$  and  $\diamond_{\beta_2}$  be two  $\beta$ -adjustment operators. Let  $\kappa_A, \kappa_B, \kappa_C$  be the same as in the proof to Proposition 7, with the added restriction that  $\kappa_A(\neg p \wedge \neg r) > \beta_1$ . As usual,  $(\kappa_A \diamond_{\beta_1} (\kappa_B \diamond_{\beta_2} \kappa_C \downarrow)) \downarrow = \|\neg p \wedge \neg r\|$ . By Definition 14,

$$\begin{aligned} & (\kappa_A \diamond_{\beta_1} B \downarrow)(p \wedge \neg r) \\ &= \beta_1 \\ &< (\kappa_A \diamond_{\beta_1} B \downarrow)(p \wedge \neg r) \\ &= \kappa_A(\neg p \wedge \neg r) \end{aligned}$$

so  $((\kappa_A \diamond_{\beta_1} \kappa_B \downarrow) \diamond_{\beta_2} \kappa_C \downarrow) \downarrow = \|p \wedge \neg r\|$ . ■

---

<sup>8</sup>Also see [21, p. 364] for a similar definition.

**Proposition 10**

$$(\Psi_1 \oslash \Psi_2)(w_1, w_2) = \Psi_1(w_1, w_2) \cap \Psi_2(w_1, w_2).$$

**Proof:** Suppose  $w_1, w_2 \in \mathcal{W}$ , and  $\Psi_1$  and  $\Psi_2$  are induced by sets of sources  $S_1, S_2 \in \mathcal{S}$ , respectively. Suppose  $s \in (\Psi_1 \oslash \Psi_2)(w_1, w_2)$ . Then, by Definitions 2 and 7,

$$\begin{aligned} s &\in \{(\mathcal{W}, \leq) \in S_1 \cap S_2 : w_1 < w_2\} \cup \{s_0\} \\ &= (\{(\mathcal{W}, \leq) \in S_1 : w_1 < w_2\} \cap \\ &\quad \{(\mathcal{W}, \leq) \in S_2 : w_1 < w_2\}) \cup \{s_0\} \\ &= (\{(\mathcal{W}, \leq) \in S_1 : w_1 < w_2\} \cup \{s_0\}) \cap \\ &\quad (\{(\mathcal{W}, \leq) \in S_2 : w_1 < w_2\} \cup \{s_0\}) \\ &= \Psi_1(w_1, w_2) \cap \Psi_2(w_1, w_2) \end{aligned}$$

Now suppose  $s \in \Psi_1(w_1, w_2) \cap \Psi_2(w_1, w_2)$ . Then, again applying Definitions 2 and 5,

$$\begin{aligned} s &\in (\{(\mathcal{W}, \leq) \in S_1 : w_1 < w_2\} \cup \{s_0\}) \cap \\ &\quad (\{(\mathcal{W}, \leq) \in S_2 : w_1 < w_2\} \cup \{s_0\}) \\ &= (\{(\mathcal{W}, \leq) \in S_1 : w_1 < w_2\} \cap \\ &\quad \{(\mathcal{W}, \leq) \in S_2 : w_1 < w_2\}) \cup \{s_0\} \\ &= \{(\mathcal{W}, \leq) \in S_1 \cap S_2 : w_1 < w_2\} \cup \{s_0\} \\ &= (\Psi_1 \oslash \Psi_2)(w_1, w_2) \end{aligned}$$

■

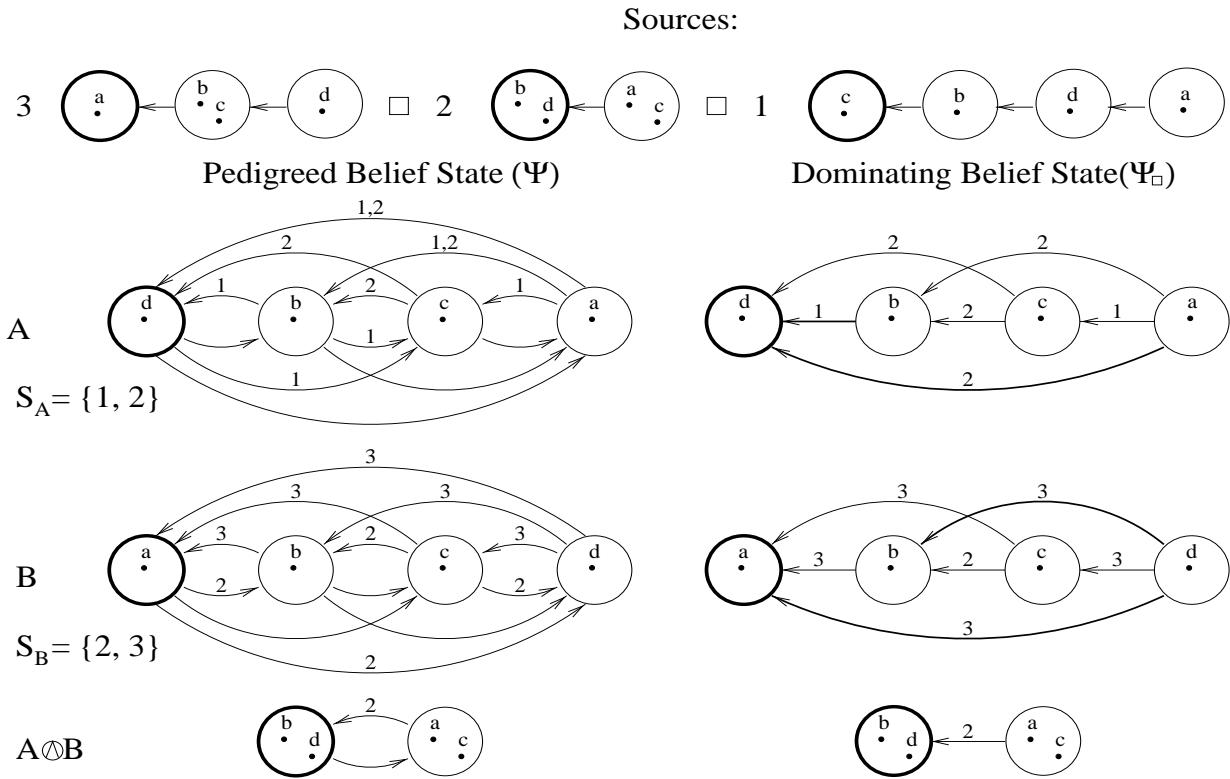


Figure 4: The diffusion operator.  $S_A$  and  $S_B$  are the sets of sources that induce the pedigreed belief states for agents  $A$  and  $B$ , respectively.

In Proc. 7th Intl. Conf. on Principles of Knowledge Representation and Reasoning (KR 2000), pp. 153–164.

# Representing and Aggregating Conflicting Beliefs

**Pedrito Maynard-Reid II**

Department of Computer Science  
Stanford University  
Stanford, CA 94305, USA  
[pedmayn@cs.stanford.edu](mailto:pedmayn@cs.stanford.edu)

**Daniel Lehmann**

School of Computer Science and Engineering  
Hebrew University  
Jerusalem 91904, Israel  
[lehmann@cs.huji.ac.il](mailto:lehmann@cs.huji.ac.il)

## Abstract

We consider the two-fold problem of representing collective beliefs and aggregating these beliefs. We propose modular, transitive relations for collective beliefs. They allow us to represent conflicting opinions and they have a clear semantics. We compare them with the quasi-transitive relations often used in Social Choice. Then, we describe a way to construct the belief state of an agent informed by a set of sources of varying degrees of reliability. This construction circumvents Arrow's Impossibility Theorem in a satisfactory manner. Finally, we give a simple set-theory-based operator for combining the information of multiple agents. We show that this operator satisfies the desirable invariants of idempotence, commutativity, and associativity, and, thus, is well-behaved when iterated, and we describe a computationally effective way of computing the resulting belief state.

**Keywords:** representation of beliefs, multi-agent systems

## 1 Introduction

We are interested in the multi-agent setting where agents are informed by sources of varying levels of reliability, and where agents can iteratively combine their belief states. This setting introduces three problems: (1) Finding an appropriate representation for collective beliefs; (2) Constructing an agent's belief state by aggregating the information from informant sources, accounting for the relative reliability of these sources; and, (3) Combining the information of multiple agents in a manner that is well-behaved under iteration.

The Social Choice community has dealt extensively with the first problem (although in the context of representing collective preferences rather than beliefs) (cf. (Sen 1986)). The classical approach has been to use quasi-transitive relations (of which total pre-orders are a special subclass) over the set of possible worlds. However, these relations do not distinguish between group indifference and group conflict, and this distinction can be crucial. Consider, for example, a situation in which all members of a group are indifferent between movie *a* and movie *b*. If some passerby expresses a preference for *a*, the group may very well choose to adopt this opinion for the group and borrow *a*. However, if the group was already divided over the relative merits of *a* and *b*, we would be wise to hesitate before choosing one over the other just because a new supporter of *a* appears on the scene. We propose a representation in which the distinction is explicit. We also argue that our representation solves some of the unpleasant semantical problems suffered by the earlier approach.

The second problem addresses how an agent should actually go about combining the information received from a set of sources to create a belief state. Such a mechanism should favor the opinions held by more reliable sources, yet allow less reliable sources to voice opinions when higher ranked sources have no opinion. True, under some circumstances it would not be advisable for an opinion from a less reliable source to override the agnosticism of a more reliable source, but often it is better to accept these opinions as default assumptions until better information is available. (Maynard-Reid II and Shoham 2000) provides a solution to this problem when belief states are represented as total pre-orders, but runs into Arrow's Impossibility Theorem (Arrow 1963) when there are sources of equal reliability. As we shall see, the generalized representation allows us to circumvent this limitation.

To motivate the third problem, consider the follow-

ing dynamic scenario: A robot controlling a ship in space receives from a number of communication centers on Earth information about the status of its environment and tasks. Each center receives information from a group of sources of varying credibility or accuracy (e.g., nearby satellites and experts) and aggregates it. Timeliness of decision-making in space is often crucial, so we do not want the robot to have to wait while each center sends its information to some central location for it to be first combined before being forwarded to the robot. Instead, each center sends its aggregated information directly to the robot. Not only does this scheme reduce dead time, it also allows for “anytime” behavior on the robot’s part: the robot incorporates new information as it arrives and makes the best decisions it can with whatever information it has at any given point. This distributed approach is also more robust since the degradation in performance is much more graceful should information from individual centers get lost or delayed.

In such a scenario, the robot needs a mechanism for combining or *fusing* the belief states of multiple agents potentially arriving at different times. Moreover, the belief state output by the mechanism should be invariant with respect to the order of agent arrivals. We will describe such a mechanism.

The paper is organized as follows: After some preliminary definitions and a discussion of the approach to aggregation taken in classical Social Choice, we introduce modular, transitive relations for representing generalized belief states. We then describe how to construct the belief state of an agent given the belief states of its informant sources when these sources are totally pre-ordered. Finally, we describe a simple set-theory-based operator for fusing agent belief states that satisfies the desirable invariants of idempotence, commutativity, and associativity, and we describe a computationally effective way of computing this belief state.

## 2 Preliminaries

We begin by defining various well-known properties of binary relations<sup>1</sup>; they will be useful to us throughout the paper.

**Definition 1** Suppose  $\leq$  is a relation over a finite set  $\Omega$ , i.e.,  $\leq \subseteq \Omega \times \Omega$ . We shall use  $x \leq y$  to denote  $(x, y) \in \leq$  and  $x \not\leq y$  to denote  $(x, y) \notin \leq$ . The relation  $\leq$  is:

---

<sup>1</sup>We only use binary relations in this paper, so we will refer to them simply as relations.

1. reflexive iff  $x \leq x$  for  $x \in \Omega$ . It is irreflexive iff  $x \not\leq x$  for  $x \in \Omega$ .
2. symmetric iff  $x \leq y \Rightarrow y \leq x$  for  $x, y \in \Omega$ . It is asymmetric iff  $x \leq y \Rightarrow y \not\leq x$  for  $x, y \in \Omega$ . It is anti-symmetric iff  $x \leq y \wedge y \leq x \Rightarrow x = y$  for  $x, y \in \Omega$ .
3. the strict version of a relation  $\leq'$  over  $\Omega$  iff  $x \leq y \Leftrightarrow x \leq' y \wedge y \not\leq' x$  for  $x, y \in \Omega$ .
4. total iff  $x \leq y \vee y \leq x$  for  $x, y \in \Omega$ .
5. modular iff  $x \leq y \Rightarrow x \leq z \vee z \leq y$  for  $x, y, z \in \Omega$ .
6. transitive iff  $x \leq y \wedge y \leq z \Rightarrow x \leq z$  for  $x, y, z \in \Omega$ .
7. quasi-transitive iff its strict version is transitive.
8. the transitive closure of a relation  $\leq'$  over  $\Omega$  iff  $x \leq y \Leftrightarrow \exists w_0, \dots, w_n \in \Omega. x = w_0 \leq' \dots \leq' w_n = y$  for some integer  $n$ , for  $x, y \in \Omega$ .
9. acyclic iff  $\forall w_0, \dots, w_n \in \Omega. w_0 < \dots < w_n$  implies  $w_n \not\leq w_0$  for all integers  $n$ , where  $<$  is the strict version of  $\leq$ .
10. a total pre-order iff it is total and transitive. It is a total order iff it is also anti-symmetric.
11. an equivalence relation iff it is reflexive, symmetric, and transitive.

### Proposition 1

1. The transitive closure of a modular relation is modular.<sup>2</sup>
2. Every transitive relation is quasi-transitive.
3. (Sen 1986) Every quasi-transitive relation is acyclic.

Given a relation over a set of alternatives and a subset of these alternatives, we often want to pick the subset’s “best” elements with respect to the relation. We define this set of “best” elements to be the subset’s *choice set*:

**Definition 2** If  $\leq$  is a relation over a finite set  $\Omega$ ,  $<$  is its strict version, and  $X \subseteq \Omega$ , then the choice set of  $X$  with respect to  $\leq$  is

$$C(X, \leq) = \{x \in X : \nexists x' \in X. x' < x\}.$$

---

<sup>2</sup>Due to space considerations, we have omitted all proofs from this manuscript. They can be found at the website <http://robotics.stanford.edu/~pedmayn/Papers/>.

A *choice function* is one which assigns to every subset  $X$  a non-empty subset of  $X$ :

**Definition 3** A choice function over a finite set  $\Omega$  is a function  $f : 2^\Omega \setminus \emptyset \rightarrow 2^\Omega \setminus \emptyset$  such that  $f(X) \subseteq X$  for every  $X \subseteq \Omega$ .

Now, every acyclic relation defines a choice function, one which assigns to each subset its choice set:

**Proposition 2** (Sen 1986) Given a relation  $\leq$  over a finite set  $\Omega$ , the choice set operation  $C$  defines a choice function iff  $\leq$  is acyclic.<sup>3</sup>

If a relation is not acyclic, elements involved in a cycle are said to be in a *conflict* because we cannot order them:

**Definition 4** Given a relation  $<$  over a finite set  $\Omega$ ,  $x$  and  $y$  are in a conflict wrt  $<$  iff there exist  $w_0, \dots, w_n, z_0, \dots, z_m \in \Omega$  such that  $x = w_0 < \dots < w_n = y = z_0 < \dots < z_m = x$ , where  $x, y \in \Omega$ .

### 3 Aggregation in Social Choice

We are interested in belief aggregation, but the community historically most interested in aggregation has been that of Social Choice theory. The aggregation is over preferences rather than beliefs, so the discussion in this subsection will focus on representing preferences; however, as we shall see, the results are equally relevant to representing beliefs. In the Social Choice community, the standard representation of an agent's preferences is a total pre-order. Each total pre-order  $\preceq_i$  is interpreted as describing the weak preferences of an individual  $i$ , so that  $x \preceq_i y$  means  $i$  considers alternative  $x$  to be at least as preferable as alternative  $y$ .<sup>4</sup> If  $x \preceq_i y$  and  $y \preceq_i x$ , then  $i$  is indifferent between  $x$  and  $y$ .

Unfortunately, Arrow's Impossibility Theorem (Arrow 1963) showed that no aggregation operator over total pre-orders exists satisfying the following small set of desirable properties:

**Definition 5** Let  $f$  be an aggregation operator over the preferences  $\preceq_1, \dots, \preceq_n$  of  $n$  individuals, respectively, over a finite set of alternatives  $\Omega$ , and let  $\preceq = f(\preceq_1, \dots, \preceq_n)$ .

---

<sup>3</sup>Sen's uses a slightly stronger definition of choice sets, but the theorem still holds in our more general case.

<sup>4</sup>The direction of the relation symbol is unintuitive, but standard practice in the belief revision community.

- Restricted Range: *The range of  $f$  is the set of total pre-orders over  $\Omega$ .*
- Unrestricted Domain: *The domain of  $f$  is the set of  $n$ -tuples of total pre-orders over  $\Omega$ .*
- Pareto Principle: *If  $x \prec_i y$  for all  $i$ , then  $x \prec y$ .*
- Independence of Irrelevant Alternatives (IIA): *Suppose  $\preceq' = f(\preceq'_1, \dots, \preceq'_n)$ . If, for  $x, y \in \Omega$ ,  $x \preceq_i y$  iff  $x \preceq'_i y$  for all  $i$ , then  $x \preceq y$  iff  $x \preceq' y$ .*
- Non-Dictatorship: *There is no individual  $i$  such that, for every tuple in the domain of  $f$  and every  $x, y \in \Omega$ ,  $x \prec_i y$  implies  $x \prec y$ .*

**Proposition 3** (Arrow 1963) There is no aggregation operator that satisfies restricted range, unrestricted domain, (weak) Pareto principle, independence of irrelevant alternatives, and nondictatorship.

This impossibility theorem led researchers to look for weakenings to Arrow's framework that would circumvent the result. One was to weaken the restricted range condition, requiring that the result of an aggregation only satisfy totality and quasi-transitivity rather than the full transitivity of a total pre-order. This weakening was sufficient to guarantee the existence of an aggregation function satisfying the other conditions, while still producing relations that defined choice functions (Sen 1986). However, this solution was not without its own problems.

First, total, quasi-transitive relations have unsatisfactory semantics. If  $\preceq$  is total and quasi-transitive but not a total pre-order, its indifference relation is not transitive:

**Proposition 4** Let  $\preceq$  be a relation over a finite set  $\Omega$  and let  $\sim$  be its symmetric restriction (i.e.,  $x \sim y$  iff  $x \preceq y$  and  $y \preceq x$ ). If  $\preceq$  is total and quasi-transitive but not transitive, then  $\sim$  is not transitive.

There has been much discussion as to whether or not indifference should be transitive; in many cases one feels indifference should be transitive. If Deb enjoys plums and mangoes equally and also enjoys mangoes and peaches equally, we would conclude that she also enjoys plums and peaches equally. It seems that total quasi-transitive relations that are not total pre-orders cannot be understood easily as preference or indifference.

Since the existence of a choice function is generally sufficient for classical Social Choice problems, this issue was at least ignorable. However, in iterated aggregation, the result of the aggregation must not only be usable for making decisions, but must be interpretable as

a new preference relation that may be involved in later aggregations; consequently, it must maintain clean semantics.

Secondly, the totality assumption is excessively restrictive for representing aggregate preferences. In general, a binary relation  $\preceq$  can express four possible relationships between a pair of alternatives  $a$  and  $b$ :  $a \preceq b$  and  $b \not\preceq a$ ,  $b \preceq a$  and  $a \not\preceq b$ ,  $a \preceq b$  and  $b \preceq a$ , and  $a \not\preceq b$  and  $b \not\preceq a$ . Totality reduces this set to the first three which, under the interpretation of relations as representing weak preference, correspond to the two strict orderings of  $a$  and  $b$ , and indifference. However, consider the situation where a couple is trying to choose between an Italian and an Indian restaurant, but one strictly prefers Italian food to Indian food, whereas the second strictly prefers Indian to Italian. The couple's opinions are in conflict, a situation that does not fit into any of the three remaining categories. Thus, the totality assumption is essentially an assumption that conflicts do not exist. This, one may argue, is appropriate if we want to represent preferences of one agent (but see (Kahneman and Tversky 1979) for persuasive arguments that individuals *are* often ambivalent). However, the assumption is inappropriate if we want to represent aggregate preferences since individuals will almost certainly have differences of opinion.

## 4 Generalized Belief States

Let us turn to the domain of belief aggregation. A total pre-order over the set of possible worlds is a fairly well-accepted representation for a belief state in the belief revision community (Grove 1988; Katsumo and Mendelzon 1991; Lehmann and Magidor 1992; Gärdenfors and Makinson 1994). Instead of preference, relations represent relative likelihood, instead of indifference, equal likelihood. For the remainder of the paper, assume we are given some language  $\mathcal{L}$  with a satisfaction relation  $\models$  for  $\mathcal{L}$ . Let  $\mathcal{W}$  be a finite, non-empty set of possible worlds (interpretations) over  $\mathcal{L}$ . Suppose  $\preceq$  is a total pre-order on  $\mathcal{W}$ . The belief revision literature maintains that the conditional belief "if  $p$  then  $q$ " (where  $p$  and  $q$  are sentences in  $\mathcal{L}$ ) holds if all the worlds in the choice set of those satisfying  $p$  also satisfy  $q$ ; we write  $Bel(p?q)$ . The individual's unconditional beliefs are all those where  $p$  is the sentence *true*. If neither the belief  $p?q$  nor its negation hold in the belief state, it is said to be *agnostic* with respect to  $p?q$ , written  $Agn(p?q)$ .

It should come as no surprise that belief aggregation is formally similar to preference aggregation and, as a result, is also susceptible to the problems described in the previous section. We propose a solution to these

problems which generalizes the total pre-order representation so as to capture information about conflicts.

### 4.1 Modular, transitive states

We take strict likelihood as primitive. Since strict likelihood is not necessarily total, it is possible to represent agnosticism and conflicting opinions in the same structure. This choice deviates from that of most authors, but are similar to those of Kreps (Kreps 1990, p. 19) who is interested in representing both indifference and incomparability. Unlike Kreps, rather than use an asymmetric relation to represent strict likelihood (e.g., the strict version of a weak likelihood relation), we impose the less restrictive condition of modularity.

We formally define *generalized belief states*:

**Definition 6** A generalized belief state  $\prec$  is a modular, transitive relation over  $\mathcal{W}$ . The set of possible generalized belief states over  $\mathcal{W}$  is denoted  $\mathcal{B}$ .

We interpret  $a \prec b$  to mean "there is reason to consider  $a$  as strictly more likely than  $b$ ." We represent equal likelihood, which we also refer to as "agnosticism," with the relationship  $\sim$  defined such that  $x \sim y$  if and only if  $x \not\prec y$  and  $y \not\prec x$ . We define the conflict relation corresponding to  $\prec$ , denoted  $\infty$ , so that  $x \infty y$  iff  $x \prec y$  and  $y \prec x$ . It describes situations where there are reasons to consider either of a pair of worlds as strictly more likely than the other. In fact, one can easily check that  $\infty$  precisely represents conflicts in a belief state in the sense of Definition 4.

For convenience, we will refer to generalized belief states simply as belief states for the remainder of the paper except when to do so would cause confusion.

### 4.2 Discussion

Let us consider why our choice of representation is justified. First, we agree with the Social Choice community that strict likelihood should be transitive.

As we discussed in the previous section, there is often no compelling reason why agnosticism/indifference should not be transitive; we also adopt this view. However, transitivity of strict likelihood by itself does not guarantee transitivity of agnosticism. A simple example is the following:  $\prec = \{(a, c)\}$ , so that  $\sim = \{(a, b), (b, c)\}$ . However, if we buy that strict likelihood should be transitive, then agnosticism is transitive identically when strict likelihood is also modular:

**Proposition 5** Suppose a relation  $\prec$  is transitive and

$\sim$  is the corresponding agnosticism relation. Then  $\sim$  is transitive iff  $\prec$  is modular.

In summary, transitivity and modularity are necessary if strict likelihood and agnosticism are both required to be transitive.

We should point out that conflicts are also transitive in our framework. At first glance, this may appear undesirable: it is entirely possible for a group to disagree on the relative likelihood of worlds  $a$  and  $b$ , and  $b$  and  $c$ , yet agree that  $a$  is more likely than  $c$ . However, we note that this transitivity follows from the cycle-based definition of conflicts (Definition 4), not from our belief state representation. It highlights the fact that we are not only concerned with conflicts that arise from simple disagreements over pairs of alternatives, but those that can be inferred from a series of inconsistent opinions as well.

Now, to argue that modular, transitive relations are sufficient to capture relative likelihood, agnosticism, and conflicts among a group of information sources, we first point out that adding irreflexivity would give us the class of relations that are strict versions of total pre-orders, i.e., conflict-free. Let  $\mathcal{T}$  be the set of total pre-orders over  $\mathcal{W}$ ,  $\mathcal{T}_<$ , the set of their strict versions.

**Proposition 6** *The set of irreflexive relations in  $\mathcal{B}$  is isomorphic to  $\mathcal{T}$  and, in fact, equals  $\mathcal{T}_<$ .*

Secondly, the following representation theorem shows that each belief state partitions the possible worlds into sets of worlds either all equally likely or all potentially involved in a conflict, and totally orders these sets; worlds in distinct sets have the same relation to each other as do the sets.

**Proposition 7**  $\prec \in \mathcal{B}$  iff there is a partition  $\mathbf{W} = \langle W_0, \dots, W_n \rangle$  of  $\mathcal{W}$  such that:

1. For every  $x \in W_i$  and  $y \in W_j$ ,  $i \neq j$  implies  $i < j$  iff  $x \prec y$ .
2. Every  $W_i$  is either fully connected ( $w \prec w'$  for all  $w, w' \in W_i$ ) or fully disconnected ( $w \not\prec w'$  for all  $w, w' \in W_i$ ).

Figure 1 shows three examples of belief states: one which is a total pre-order, one which is the strict version of a total pre-order, and one which is neither.

Thus, generalized belief states are not a big change from the strict versions of total pre-orders. They merely generalize these by weakening the assumption that sets of worlds not strictly ordered are equally likely, allowing for the possibility of conflicts. Now

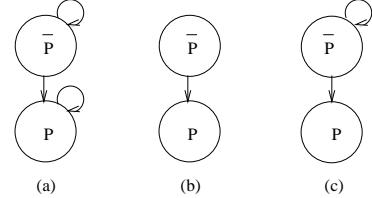


Figure 1: Three examples of generalized belief states: (a) a total pre-order, (b) the strict version of a total pre-order, (c) neither. (Each circle represents all the worlds in  $\mathcal{W}$  which satisfy the sentence inside. An arc between circles indicates that  $w \prec w'$  for every  $w$  in the head circle and  $w'$  in the tail circle; no arc indicates that  $w \not\prec w'$  for each of these pairs. In particular, the set of worlds represented by a circle is fully connected if there is an arc from the circle to itself, fully disconnected otherwise.)

we can distinguish between agnostic and conflicting conditional beliefs. A belief state  $\prec$  is agnostic about conditional belief  $p?q$  (i.e.,  $Agn(p?q)$ ) if the choice set of worlds satisfying  $p$  contains both worlds which satisfy  $q$  and  $\neg q$  and is fully disconnected. It is in conflict about this belief, written  $Con(p?q)$ , if the choice set is fully connected.

Finally, we compare the representational power of our definitions to those discussed in the previous section. First,  $\mathcal{B}$  subsumes the class of total pre-orders:

**Proposition 8**  $\mathcal{T} \subset \mathcal{B}$  and is the set of reflexive relations in  $\mathcal{B}$ .

Secondly,  $\mathcal{B}$  neither subsumes nor is subsumed by the set of total, quasi-transitive relations, and the intersection of the two classes is  $\mathcal{T}$ . Let  $\mathcal{Q}$  be the set of total, quasi-transitive relations over  $\mathcal{W}$ , and  $\mathcal{Q}_<$ , the set of their strict versions.

**Proposition 9**

1.  $\mathcal{Q} \cap \mathcal{B} = \mathcal{T}$ .
2.  $\mathcal{B} \not\subseteq \mathcal{Q}$ .
3.  $\mathcal{Q} \not\subseteq \mathcal{B}$  if  $\mathcal{W}$  has at least three elements.
4.  $\mathcal{Q} \subset \mathcal{B}$  if  $\mathcal{W}$  has one or two elements.

Because modular, transitive relations represent strict preferences, it is probably fairer to compare them to the class of strict versions of total, quasi-transitive relations. Again, neither class subsumes the other, but this time the intersection is  $\mathcal{T}_<$ :

### Proposition 10

1.  $\mathcal{Q}_< \cap \mathcal{B} = \mathcal{T}_<$ .
2.  $\mathcal{B} \not\subseteq \mathcal{Q}_<$ .
3.  $\mathcal{Q}_< \not\subseteq \mathcal{B}$  if  $\mathcal{W}$  has at least three elements.
4.  $\mathcal{Q}_< \subset \mathcal{B}$  if  $\mathcal{W}$  has one or two elements.

In the next section, we define a natural aggregation policy based on this new representation that admits clear semantics and obeys appropriately modified versions of Arrow's conditions.

## 5 Single-agent belief state construction

Suppose an agent is informed by a set of sources, each with its individual belief state. Suppose further that the agent has ranked the sources by level of credibility. We propose an operator for constructing the agent's belief state  $\prec$  by aggregating the belief states of the sources in  $S$  while accounting for the credibility ranking of the sources.

**Example 1** *We will use a running example from our space robot domain to help provide intuition for our definitions. The robot sends to earth a stream of telemetry data gathered by the spacecraft, as long as it receives positive feedback that the data is being received. At some point it loses contact with the automatic feedback system, so it sends a request for information to an agent on earth to find out if the failure was caused by a failure of the feedback system or by an overload of the data retrieval system. In the former case, it would continue to send data, in the latter, desist. As it so happens, there has been no overload, but the computer running the feedback system has hung. The agent consults the following three experts, aggregates their beliefs, and sends the results back to the robot:*

1.  $s_p$ , the computer programmer that developed the feedback program, believes nothing could ever go wrong with her code, so there must have been an overload problem. However, she admits that if her program had crashed, the problem could ripple through to cause an overload.
2.  $s_m$ , the manager for the telemetry division, unfortunately has out-dated information that the feedback system is working. She was also told by the engineer who sold her the system that overloading could never happen. She has no idea what would

happen if there was an overload or the feedback system crashed.

3.  $s_t$ , the technician working on the feedback system, knows that the feedback system crashed, but doesn't know whether there was a data-overload. Not being familiar with the retrieval system, she is also unable to speculate whether the data retrieval system would have overloaded if the feedback system had not failed.

Let  $F$  and  $D$  be propositional variables representing that the feedback and data retrieval systems, respectively, are okay. The belief states for the three sources are shown in Figure 2.

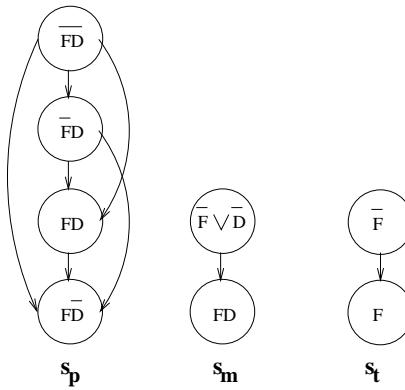


Figure 2: The belief states of  $s_p$ ,  $s_m$ , and  $s_t$  in Example 1.

Let us begin the formal development by defining sources:

**Definition 7**  $\mathcal{S}$  is a finite set of sources. With each source  $s \in \mathcal{S}$  is associated a belief state  $\prec^s \in \mathcal{B}$ .

We denote the agnosticism and conflict relations of a source  $s$  by  $\approx^s$  and  $\bowtie^s$ , respectively. It is possible to assume that the belief state of a source is conflict free, i.e., acyclic. However, this is not necessary if we allow sources to suffer from the human malady of "being torn between possibilities."

We assume that the agent's credibility ranking over the sources is a total pre-order:

**Definition 8**  $\mathcal{R}$  is a totally ordered finite set of ranks.

**Definition 9**  $\text{rank} : \mathcal{S} \rightarrow \mathcal{R}$  assigns to each source a rank.

**Definition 10**  $\sqsupseteq$  is the total pre-order over  $\mathcal{S}$  induced by the ordering over  $\mathcal{R}$ . That is,  $s \sqsupseteq s'$  iff

$\text{rank}(s) \geq \text{rank}(s')$ ; we say  $s'$  is as credible as  $s$ .  $\sqsupseteq_S$  is the restriction of  $\sqsupseteq$  to  $S \subseteq \mathcal{S}$ .

We use  $\sqsupset$  and  $\equiv$  to denote the asymmetric and symmetric restrictions of  $\sqsupseteq$ , respectively.<sup>5</sup> The finiteness of  $\mathcal{S}$  ( $\mathcal{R}$ ) ensures that a maximal source (rank) always exists, which is necessary for some of our results. Weaker assumptions are possible, but at the price of unnecessarily complicating the discussion.

We are ready to consider the source aggregation problem. In the following, assume an agent is informed by a set of sources  $S \subseteq \mathcal{S}$ . We look at two special cases—equal-ranked and strictly-ranked source aggregation—before considering the general case.

### 5.1 Equal-ranked sources aggregation

Suppose all the sources have the same rank so that  $\sqsupseteq_S$  is fully connected. Intuitively, we want take all offered opinions seriously, so we take the union of the relations:

**Definition 11** If  $S \subseteq \mathcal{S}$ , then  $\text{Un}(S)$  is the relation  $\bigcup_{s \in S} <^s$ .

By simply taking the union of the source belief states, we may lose transitivity. However, we do not lose modularity:

**Proposition 11** If  $S \subseteq \mathcal{S}$ , then  $\text{Un}(S)$  is modular but not necessarily transitive.

Thus, we know from Proposition 1 that we need only take the transitive closure of  $\text{Un}(S)$  to get a belief state:

**Definition 12** If  $S \subseteq \mathcal{S}$ , then  $\text{AGRUn}(S)$  is the relation  $\text{Un}(S)^+$ .

**Proposition 12** If  $S \subseteq \mathcal{S}$ , then  $\text{AGRUn}(S) \in \mathcal{B}$ .

Not surprisingly, by taking all opinions of all sources seriously, we may generate many conflicts, manifested as fully connected subsets of  $\mathcal{W}$ .

**Example 2** Suppose all three sources in the space robot scenario of Example 1 are considered equally credible, then the aggregate belief state will be the fully connected relation indicating that there are conflicts over every belief.

<sup>5</sup>Note that, unlike the relations representing belief states,  $\geq$  and  $\sqsupseteq$  are read in the intuitive way, that is, “greater” corresponds to “better.”

### 5.2 Strictly-ranked sources aggregation

Next, consider the case where the sources are strictly ranked, i.e.,  $\sqsupseteq_S$  is a total order. We define an operator such that lower-ranked sources refine the belief states of higher ranked sources. That is, in determining the ordering of a pair of worlds, the opinions of higher-ranked sources generally override those of lower-ranked sources, and lower-ranked sources are consulted when higher-ranked sources are agnostic:

**Definition 13** If  $S \subseteq \mathcal{S}$ , then  $\text{AGRf}(S)$  is the relation

$$\{(x, y) : \exists s \in S. x <^s y \wedge (\forall s' \sqsupset s \in S. x \approx^{s'} y)\}.$$

The definition of the  $\text{AGRf}$  operator does not rely on  $\sqsupseteq_S$  being a total order, and we will use it in this more general setting in the following sub-section. However, in the case that  $\sqsupseteq_S$  is a total order, the result of applying  $\text{AGRf}$  is guaranteed to be a belief state.

**Proposition 13** If  $S \subseteq \mathcal{S}$  and  $\sqsupseteq_S$  is a total order, then  $\text{AGRf}(S) \in \mathcal{B}$ .

**Example 3** Suppose, in the space robot scenario of Example 1, the technician is considered more credible than the manager who, in turn, is considered more credible than the programmer. The aggregate belief state, shown in Figure 3, informs the robot correctly that the feedback system has crashed, but that it shouldn’t worry about an overload problem and should keep sending data.

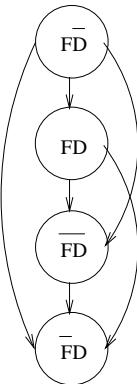


Figure 3: The belief state after aggregation in Example 3 when  $s_t \sqsupset s_m \sqsupset s_p$ .

Note that this case of strictly-ranked sources is almost exactly that considered in (Maynard-Reid II and Shoham 2000), except that the authors are not able to allow for conflicts in belief states. A surprising result they show is that standard AGM belief revision (Alchourrón *et al.* 1985) can be modeled as the aggrega-

tion of two sources, the informant and the informee, where the informant is considered more credible than the informee.

### 5.3 General aggregation

In the general case, we may have several ranks represented and multiple sources of each rank. It will be instructive to first consider the following seemingly natural strawman operator,  $AGR^*$ : First combine equi-rank sources using  $AGRUn$ , then aggregate the strictly-ranked results using what is essentially  $AGRRf$ :

**Definition 14** Let  $S \subseteq \mathcal{S}$ . For any  $r \in \mathcal{R}$ , let  $<_r = AGRUn(\{s \in S : \text{rank}(s) = r\})$  and  $\approx_{r'}$ , the corresponding agnosticism relation. Also, let  $\text{rank}(S) = \{r \in \mathcal{R} : \exists s \in S. \text{rank}(s) = r\}$ .  $AGR^*(S)$  is the relation

$$\left\{ (x, y) : \begin{array}{l} \exists r \in \mathcal{R}. x <_r y \wedge \\ (\forall r' > r \in \text{rank}(S). x \approx_{r'} y) \end{array} \right\}$$

$AGR^*$  indeed defines a legitimate belief state:

**Proposition 14** If  $S \subseteq \mathcal{S}$ , then  $AGR^*(S) \in \mathcal{B}$ .

Unfortunately, a problem with this “divide-and-conquer” approach is it assumes the result of aggregation is independent of potential interactions between the individual sources of different ranks. Consequently, opinions that will eventually get overridden may still have an indirect effect on the final aggregation result by introducing superfluous opinions during the intermediate equi-rank aggregation step, as the following example shows:

**Example 4** Let  $\mathcal{W} = \{a, b, c\}$ . Suppose  $S \subseteq \mathcal{S}$  such that  $S = \{s_0, s_1, s_2\}$  with belief states  $<^{s_0} = \{(b, a), (b, c)\}$  and  $<^{s_1} = <^{s_2} = \{(a, b), (c, b)\}$ , and where  $s_2 \sqsupset s_1 \equiv s_0$ . Then  $AGR^*(S)$  is  $\{(a, b), (c, b), (a, c), (c, a), (a, a), (b, b), (c, c)\}$ . All sources are agnostic over  $a$  and  $c$ , yet  $(a, c)$  and  $(c, a)$  are in the result because of the transitive closure in the lower rank involving opinions  $((b, c))$  and  $((b, a))$  which actually get overridden in the final result.

Because of these undesired effects, we propose another aggregation operator which circumvents this problem by applying refinement (as defined in Definition 13) to the set of source belief states before inferring new opinions via closure:

**Definition 15** The rank-based aggregation of a set of sources  $S \subseteq \mathcal{S}$  is  $AGR(S) = AGRRf(S)^+$ .

Encouragingly,  $AGR$  outputs a valid belief state:

**Proposition 15** If  $S \subseteq \mathcal{S}$ , then  $AGR(S) \in \mathcal{B}$ .

**Example 5** Suppose, in the space robot scenario of Example 1, the technician is still considered more credible than the manager and the programmer, but the latter two are considered equally credible. The aggregate belief state, shown in Figure 5, still gives the robot the correct information about the state of the system. The robot also learns for future reference that there is some disagreement over whether or not there would have been a data overload if the feedback system were working.

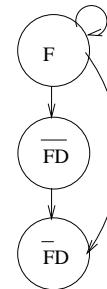


Figure 4: The belief state after aggregation in Example 5 when  $s_t \sqsupset s_m \equiv s_p$ .

We observe that  $AGR$ , when applied to the set of sources in Example 4, does indeed bypass the problem described above of extraneous opinion introduction:

**Example 6** Assume  $\mathcal{W}$ ,  $S$ , and  $\sqsupset$  are as in Example 4.  $AGR(S) = \{(a, b), (c, b)\}$ .

We also observe that  $AGR$  behaves well in the special cases we’ve considered, reducing to  $AGRUn$  when all sources have equal rank, and to  $AGRRf$  when the sources are totally ranked:

**Proposition 16** Suppose  $S \subseteq \mathcal{S}$ .

1. If  $\sqsupset_S$  is fully connected,  $AGR(S) = AGRUn(S)$ .
2. If  $\sqsupset_S$  is a total order,  $AGR(S) = AGRRf(S)$ .

### 5.4 Arrow, revisited

Finally, a strong argument in favor of  $AGR$  is that it satisfies appropriate modifications of Arrow’s conditions. Let  $f$  be an operator which aggregates the belief states  $<^{s_1}, \dots, <^{s_n}$  over  $\mathcal{W}$  of  $n$  sources  $s_1, \dots, s_n \in \mathcal{S}$ , respectively, and let  $\prec = f(<^{s_1}, \dots, <^{s_n})$ . We consider each condition separately.

**Restricted range** The output of the aggregation function will be a modular, transitive belief state rather than a total pre-order.

**Definition 16** (modified) Restricted Range: *The range of  $f$  is  $\mathcal{B}$ .*

**Unrestricted domain** Similarly, the input to the aggregation function will be modular, transitive belief states of sources rather than total pre-orders.

**Definition 17** (modified) Unrestricted Domain: *For each  $i$ ,  $<^{s_i}$  can be any member of  $\mathcal{B}$ .*

**Pareto principle** Generalized belief states already represent strict likelihood. Consequently, we use the actual input and output relations of the aggregation function in place of their strict versions to define the Pareto principle. Obviously, because we allow for the introduction of conflicts,  $AGR$  will not satisfy the original formal Pareto principle which essentially states that if all sources have an unconflicted belief that one world is strictly more likely than another, this must also be true of the aggregated belief state. Neither condition is necessarily stronger than the other.

**Definition 18** (modified) Pareto Principle: *If  $x <^{s_i} y$  for all  $i$ , then  $x \prec y$ .*

**Independence of irrelevant alternatives** Conflicts are defined in terms of cycles, not necessarily binary. By allowing the existence of conflicts, we effectively have made it possible for outside worlds to affect the relation between a pair of worlds, viz., by involving them in a cycle. As a result, we need to weaken IIA to say that the relation between worlds should be independent of other worlds **unless** these other worlds put them in conflict.

**Definition 19** (modified) Independence of Irrelevant Alternatives (IIA): *Suppose  $s'_1, \dots, s'_n \in \mathcal{S}$  such that  $s_i \equiv s'_i$  for all  $i$ , and  $\prec' = f(<^{s'_1}, \dots, <^{s'_n})$ . If, for  $x, y \in \mathcal{W}$ ,  $x <^{s_i} y$  iff  $x <^{s'_i} y$  for all  $i$ ,  $x \not\propto y$ , and  $x \not\propto' y$ , then  $x \prec y$  iff  $x \prec' y$ .*

**Non-dictatorship** As with the Pareto principle definition, we use the actual input and output relations to define non-dictatorship since belief states represent strict likelihood. From this perspective, our setting requires that informant sources of the highest rank be “dictators” in the sense considered by Arrow. However, the setting originally considered by Arrow was one where all individuals are ranked equally. Thus, we make this explicit in our new definition of

non-dictatorship by adding the pre-condition that all sources be of equal rank. Now,  $AGR$  treats a set of equi-rank sources equally by taking all their opinions seriously, at the price of introducing conflicts. So, intuitively, there are no dictators. However, because Arrow did not account for conflicts in his formulation, all the sources will be “dictators” by his definition. We need to modify the definition of non-dictatorship to say that no source can always push opinions through without them ever being contested.

**Definition 20** (modified) Non-Dictatorship: *If  $s_i \equiv s_j$  for all  $i, j$ , then there is no  $i$  such that, for every combination of source belief states and every  $x, y \in \mathcal{W}$ ,  $x <^{s_i} y$  and  $y \not\prec s_i x$  implies  $x \prec y$  and  $y \not\prec x$ .*

We now show that  $AGR$  indeed satisfies these conditions:

**Proposition 17** Let  $S = \{s_1, \dots, s_n\} \subseteq \mathcal{S}$  and  $AGR_f(<^{s_1}, \dots, <^{s_n}) = AGR(S)$ .  $AGR_f$  satisfies (the modified versions of) restricted range, unrestricted domain, Pareto principle, IIA, and non-dictatorship.

## 6 Multi-agent fusion

So far, we have only considered the case where a single agent must construct or update her belief state once informed by a set of sources. Multi-agent fusion is the process of aggregating the belief states of a set of agents, each with its respective set of informant sources. We proceed to formalize this setting.

An agent  $A$  is *informed* by a set of sources  $S \subseteq \mathcal{S}$ . Agent  $A$ 's *induced belief state* is the belief state formed by aggregating the belief states of its informant sources, i.e.,  $AGR(S)$ . Assume the set of agents to fuse agree upon *rank* (and, consequently,  $\sqsupseteq$ ).<sup>6</sup> We define the fusion of this set to be an agent informed by the combination of informant sources:

**Definition 21** Let  $\mathcal{A} = \{A_1, \dots, A_n\}$  be a set of agents such that each agent  $A_i$  is informed by  $S_i \subseteq \mathcal{S}$ . The fusion of  $\mathcal{A}$ , written  $\bigcircledcirc(\mathcal{A})$ , is an agent informed by  $S = \bigcup_{i=1}^n S_i$ .

---

<sup>6</sup>We could easily extend the framework to allow for individual rankings, but we felt that the small gain in generality would not justify the additional complexity and loss of perspicuity. Similarly, we could consider each agent as having a credibility ordering only over its informant sources. However, it is unclear how, for example, credibility orderings over disjoint sets of sources should be combined into a new credibility ordering since their union will not be total.

Not surprisingly given its set-theoretic definition, fusion is idempotent, commutative, and associative. These properties guarantee the invariance required in multi-agent belief aggregation applications such as our space robot domain.

In the multi-agent space robot scenario described in Section 1, we only have a direct need for the belief states that result from fusion. We are only interested in the belief states of the original sources in as far as we want the fused belief state to reflect its informant history. An obvious question is whether it is possible to compute the belief state induced by the agents' fusion solely from their initial belief states, that is, without having to reference the belief states of their informant sources. This is highly desirable because of the expense of storing—or, as in the case of our space robot example, transmitting—all source belief states; we would like to represent each agent's knowledge as compactly as possible.

In fact, we can do this if all sources have equal rank. We simply take the transitive closure of the union of the agents' belief states:

**Proposition 18** *Let  $\mathcal{A}$  and  $S$  be as in Definition 21,  $\prec^{A_i}$ , agent  $A_i$ 's induced belief state, and  $\sqsupseteq_S$ , fully connected. If  $A = \bigcirclearrowright(\mathcal{A})$ , then  $(\bigcup_{A_i \in \mathcal{A}} \prec^{A_i})^+$  is  $A$ 's induced belief state.*

Unfortunately, the equal rank case is special. If we have sources of different ranks, we generally cannot compute the induced belief state after fusion using only the agent belief states before fusion, as the following simple example demonstrates:

**Example 7** *Let  $\mathcal{W} = \{a, b\}$ . Suppose two agents  $A_1$  and  $A_2$  are informed by sources  $s_1$  with belief state  $\prec^{s_1} = \{(a, b)\}$  and  $s_2$  with belief state  $\prec^{s_2} = \{(b, a)\}$ , respectively.  $A_1$ 's belief state is the same as  $s_1$ 's and  $A_2$ 's is the same as  $s_2$ 's. If  $s_1 \sqsupseteq s_2$ , then the belief state induced by  $\bigcirclearrowright(A_1, A_2)$  is  $\prec^{s_1}$ , whereas if  $s_2 \sqsupseteq s_1$ , then it is  $\prec^{s_2}$ . Thus, just knowing the belief states of the fused agents is not sufficient for computing the induced belief state. We need more information about the original sources.*

However, if sources are totally pre-ordered by credibility, we can still do much better than storing all the original sources. It is enough to store for each opinion of  $AGRf(S)$  the rank of the highest-ranked source supporting it. We define *pedigreed belief states* which enrich belief states with this additional information:

**Definition 22** *Let  $A$  be an agent informed by a set of sources  $S \subseteq \mathcal{S}$ .  $A$ 's pedigreed belief state is a pair*

$(\prec, l)$  where  $\prec = AGRf(S)$  and  $l : \prec \rightarrow \mathcal{R}$  such that  $l((x, y)) = \max\{\text{rank}(s) : x \prec^s y, s \in S\}$ . We use  $\prec_r^A$  to denote the restriction of  $A$ 's pedigreed belief state to  $r$ , that is,  $\prec_r^A = \{(x, y) \in \prec : l((x, y)) = r\}$ .

We verify that a pair's label is, in fact, the rank of the source used to determine the pair's membership in  $AGRf(S)$ , not that of some higher-ranked source:

**Proposition 19** *Let  $A$  be an agent informed by a set of sources  $S \subseteq \mathcal{S}$  and with pedigreed belief state  $(\prec, l)$ . Then*

$$x \prec_r^A y$$

iff

$$\exists s \in S. x \prec^s y \wedge r = \text{rank}(s) \wedge (\forall s' \sqsupseteq s \in S. x \approx^{s'} y).$$

The belief state induced by a pedigreed belief state  $(\prec, l)$  is, obviously, the transitive closure of  $\prec$ .

Now, given only the pedigreed belief states of a set of agents, we can compute the new pedigreed belief state after fusion. We simply combine the labeled opinions using our refinement techniques.

**Proposition 20** *Let  $\mathcal{A}$  and  $S$  be as in Definition 21,  $\sqsupseteq_S$ , a total pre-order, and  $A = \bigcirclearrowright(\mathcal{A})$ . If*

1.  $\prec$  is the relation

$$\left\{ (x, y) : \begin{array}{l} \exists A_i \in \mathcal{A}, r \in \mathcal{R}. x \prec_r^{A_i} y \wedge \\ (\forall A_j \in \mathcal{A}, r' > r \in \mathcal{R}. x \sim_{r'}^{A_j} y) \end{array} \right\}$$

over  $\mathcal{W}$ ,

2.  $l : \prec \rightarrow \mathcal{R}$  such that  $l((x, y)) = \max\{r : x \prec_r^{A_i} y, A_i \in \mathcal{A}\}$ , and

then  $(\prec, l)$  is  $A$ 's pedigreed belief state.

From the perspective of the induced belief states, we are essentially discarding unlabeled opinions (i.e., those derived by the closure operation) before fusion. Intuitively, we are learning new information so we may need to retract some of our inferred opinions. After fusion, we re-apply closure to complete the new belief state. Interestingly, in the special case where the sources are strictly-ranked, the closure is unnecessary:

**Proposition 21** *If  $\mathcal{A}$  and  $S$  are as in Definition 21,  $\sqsupseteq_S$  is a total order, and  $(\prec, l)$  is the pedigreed belief state of  $\bigcirclearrowright(\mathcal{A})$ , then  $\prec^+ = \prec$ .*

**Example 8** Let's look once more at the space robot scenario considered in Example 1. Suppose the arrogant programmer is not part of the telemetry team, but instead works for a company on the other side of the country. Then the robot has to request information from two separate agents, one to query the manager and technician and one to query the programmer. Assume that the agents and the robot all rank the sources the same, assigning the technician rank 2 and the other two agents rank 1, which induces the same credibility ordering used in Example 5. The agents' pedigreed belief states and the result of their fusion are shown in Figure 5.

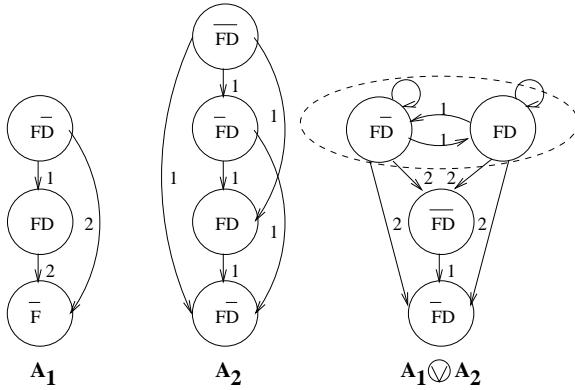


Figure 5: The pedigreed belief states of agent  $A_1$  informed by  $s_m$  and  $s_t$  and of agent  $A_2$  informed by  $s_p$ , and the result of their fusion in Example 8.

The first agent does not provide any information about overloading and the second agent provides incorrect information. However, we see that after fusing the two, the robot has a belief state that is identical to what it computed in Example 5 when there was only one agent informed by all three sources (we've only separated the top set of worlds so as to show the labeling). Consequently, it now knows the correct state of the system. And, satisfyingly, the final result does not depend on the order in which the robot receives the agents' reports.

The savings obtained in required storage space by this scheme can be substantial. Whereas explicitly storing all of an agent's informant sources  $S$  requires  $\mathcal{O}(|S|2^W)$  amount of space in the worst case (when all the sources' belief states are fully connected relations), storing a pedigreed belief state only requires  $\mathcal{O}(2^W)$  space in the worst case. Moreover, not only does the enriched representation allow us to conserve space, but it also provides for potential savings in the efficiency of computing fusion since, for each pair of worlds, we only need to consider the opinions of the

agents rather than those of all the sources in the combined set of informants.

Incidentally, if we had used  $AGR^*$  as the basis for our general aggregation, simply storing the rank of the maximum supporting sources would not give us sufficient information to compute the induced belief state after fusion. To demonstrate this, we give an example where two pairs of sources induce the same annotated agent belief states, yet yield different belief states after fusion:

**Example 9** Let  $\mathcal{W}$ ,  $\mathcal{S}$ , and  $\sqsupseteq$  be as in Example 4. Suppose agents  $A_1$ ,  $A_2$ ,  $A'_1$ , and  $A'_2$  are informed by sets of sources  $S_1$ ,  $S_2$ ,  $S'_1$ , and  $S'_2$ , respectively, where  $S_1 = S_2 = \{s_2\}$ ,  $S'_1 = \{s_0, s_2\}$ , and  $S'_2 = \{s_1, s_2\}$ .  $AGR^*$  dictates that the pedigreed belief states of all four agents equal  $<^{s_2}$  with all opinions annotated with  $\text{rank}(s_2)$ . In spite of this indistinguishability, if  $A = \bigvee(\{A_1, A_2\})$  and  $A' = \bigvee(\{A'_1, A'_2\})$ , then  $A$ 's induced belief state equals  $<^{s_2}$ , i.e.,  $\{(a, b), (c, b)\}$ , whereas  $A'$ 's is  $\{(a, b), (c, b), (a, c), (c, a), (b, b), (c, c)\}$ .

## 7 Conclusion

We have described a semantically clean representation for aggregate beliefs which allows us to represent conflicting opinions without sacrificing the ability to make decisions. We have proposed an intuitive operator which takes advantage of this representation so that an agent can combine the belief states of a set of informant sources totally pre-ordered by credibility. Finally, we have described a mechanism for fusing the belief states of different agents which iterates well.

The aggregation methods we have discussed here are just special cases of a more general framework based on voting. That is, we account not only for the ranking of the sources supporting or disagreeing with an opinion (i.e., the *quality* of support), but also the percentage of sources in each camp (the *quantity* of support). Such an extension allows for a much more refined approach to aggregation, one much closer to what humans often use in practice. Exploring this richer space is the subject of further research.

Another problem which deserves further study is developing a fuller understanding of the properties of the *Bel*, *Agn*, and *Con* operators and how they interrelate.

## Acknowledgements

Pedrito Maynard-Reid II was partly supported by a National Physical Science Consortium Fellowship. The final version of this paper was written with the

financial support of the Jean et Hélène Alfassa Fund for Research in Artificial Intelligence.

## References

- Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- Kenneth J. Arrow. *Social Choice and Individual Values*. Wiley, New York, 2nd edition, 1963.
- Peter Gärdenfors and David Makinson. Nonmonotonic inference based on expectations. *Artificial Intelligence*, 65(1):197–245, January 1994.
- Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.
- D. Kahneman and A. Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, March 1979.
- Hirofumi Katsuno and Alberto O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3):263–294, 1991.
- David M. Kreps. *A Course in Microeconomic Theory*. Princeton University Press, 1990.
- Daniel Lehmann and Menachem Magidor. What does a conditional knowledge base entail? *Artificial Intelligence*, 55(1):1–60, May 1992.
- Pedrito Maynard-Reid II and Yoav Shoham. Belief fusion: Aggregating pedigreed belief states. *Journal of Logic, Language, and Information*, 2000. To appear.
- Amartya Sen. Social choice theory. In K. J. Arrow and M. D. Intriligator, editors, *Handbook of Mathematical Economics*, volume III, chapter 22, pages 1073–1181. Elsevier Science Publishers, 1986.

---

## A Normative Examination of Ensemble Learning Algorithms

---

**David M. Pennock**

NEC Research Institute, 4 Independence Way, Princeton, NJ 08540 USA

D PENNOCK@RESEARCH.NJ.NEC.COM

**Pedrito Maynard-Reid II**

Computer Science Department, Stanford University, Stanford, CA 94305 USA

PEDMAYN@CS.STANFORD.EDU

**C. Lee Giles**

NEC Research Institute, 4 Independence Way, Princeton, NJ 08540 USA

GILES@RESEARCH.NJ.NEC.COM

**Eric Horvitz**

Microsoft Research, One Microsoft Way, Redmond, WA 98052-6399 USA

HORVITZ@MICROSOFT.COM

### Abstract

Ensemble learning algorithms combine the results of several classifiers to yield an aggregate classification. We present a normative evaluation of combination methods, applying and extending existing axiomatizations from Social Choice theory and Statistics. For the case of multiple classes, we show that several seemingly innocuous and desirable properties are mutually satisfied only by a *dictatorship*. A weaker set of properties admit only the weighted average combination rule. For the case of binary classification, we give axiomatic justifications for majority vote and for weighted majority. We also show that, even when all component algorithms report that an attribute is probabilistically independent of the classification, common ensemble algorithms often destroy this independence information. We exemplify these theoretical results with experiments on stock market data, demonstrating how ensembles of classifiers can exhibit canonical voting paradoxes.

center on generalization accuracy, either deriving theoretical bounds (Schapire, 1990; Freund & Schapire, 1999) or (more commonly) comparing experimental results (Bauer & Kohavi, 1999; Breiman, 1996; Dietterich, in press; Freund & Schapire, 1996).

We take instead a normative approach, informed by results from Social Choice theory and statistical belief aggregation. First, we identify several properties that an ensemble algorithm might ideally possess, and then characterize the implied form of the combination function. Section 4 examines the case of more than two classes. We show that, under a set of seemingly mild and reasonable conditions, *no* true combination method is possible. The aggregate classification is always identical to that of only one of the component algorithms. The analysis mirrors Arrow's celebrated Impossibility Theorem, which shows that the only voting mechanism that obeys a similar set of properties is a dictatorship (Arrow, 1963). Under slightly weaker demands, we show that the only possible form for the combination function is a weighted average of the constituent classifications.

Section 5 considers the special case of binary classification. Based on May's (1952) seminal work, we present a set of axioms that necessitate the use of simple majority vote to combine classifiers. We then extend this result, deriving an axiomatic justification for the *weighted* majority vote. Majority and weighted majority are two of the most common methods used for classifier combination (Dietterich, 1997). One contribution of this paper is to provide formal justifications for them.

Section 6 explores the independence preservation properties of common ensemble learning algorithms. Suppose that, with some attribute values missing, all of the constituent algorithms judge one attribute to be statistically independent of the classification. We demonstrate that this

## 1. Introduction

A recent trend in machine learning is to aggregate the outputs of several learning algorithms together to produce a composite classification (Dietterich, 1997). Under favorable conditions, ensemble classifiers provably outperform their constituent algorithms, an advantage born out by much empirical validation. Yet there does not seem to be a single, obvious way to combine classifiers—many different methods have been proposed and tested, with none emerging as the clear winner. Most evaluation metrics

independence is generally lost after combination, rendering the aggregate classification statistically *dependent* on the attribute in question.

Section 7 presents empirical evidence of violations of the various axioms. We show that an ensemble of neural networks—trained to predict stock market data—can generate counterintuitive results, reminiscent of so-called *voting paradoxes* in the Social Choice literature. Section 8 summarizes and discusses future work.

## 2. Ensemble Learning

We present a very brief overview of ensemble learning; see (Dietterich, 1997) for an excellent survey. Representative algorithms include bagging (Breiman, 1996), boosting (e.g., ADABOOST (Freund & Schapire, 1999)), and a method based on Error-Correcting Output Codes (ECOC) (Dietterich & Bakiri, 1995). Ensemble algorithms generally proceed in two phases: (1) generate and train a set of weak learners, and (2) aggregate their classifications.

The first step is to construct component learners of sufficient diversity (Hansen & Salamon, 1990). One common technique is to subsample the training examples, either randomly with replacement (Breiman, 1996), by leaving out random subsets (as in cross-validation), or by an induced distribution meant to magnify the effect of difficult training examples (Freund & Schapire, 1999). Another technique bases each learner’s predictions on different input features (Tumer & Ghosh, 1996). The method of Error-Correcting Output Codes (ECOC) generates classifiers by having each learn whether an example falls within a randomly chosen subset of the classes. Another approach injects randomness into the training algorithms themselves. These four techniques apply to arbitrary classifier algorithms—there are also many algorithm-specific techniques. And, of course, it is possible to create an ensemble by mixing and matching different techniques for different classifiers.

After generating and training a set of weak learners, the ensemble algorithm combines the individual learners’ predictions into a composite prediction. The choice of combination method is the focus of this paper. Common methods can be categorized loosely into two categories: those that combine *votes*, and those that can combine confidence scores. The former type includes plurality vote<sup>1</sup> and weighted plurality; the latter includes stacking, serial combination, weighted average, and weighted geometric average.

Bagging and ECOC are examples of algorithms that use plurality vote. The ensemble’s chosen class is simply that

<sup>1</sup>This is the familiar “one person, one vote” procedure where the candidate receiving the most votes wins. We reserve *majority vote* to refer to the special case of two candidates.

which is predicted most often by the individual learners. Weighted plurality is a generalization of plurality vote, where each algorithm’s vote is discounted (or magnified) by a multiplicative weight; classes are then ranked according to the sum of the weighted votes they receive. Weights can be chosen to correspond with the observed accuracy of the individual classifiers, using Bayesian techniques, or using *gating networks* (Jordan & Jacobs, 1994), among other methods. The ADABOOST algorithm computes weights in an attempt to minimize the error of the final classification.

*Stacking* turns the problem of finding a good combination function into a learning problem itself (Breiman, 1996; Lee & Srihari, 1995; Wolpert, 1992): The constituent algorithms’ outputs are fed to a meta learner’s inputs; the meta learner’s output is taken as the ensemble classification. *Serial combination* uses one learner’s top  $k$  choices to reduce the space of candidate classes, passing the simplified problem onto the next learner, etc. (Madhvanath & Govindaraju, 1995). *Weighted algebraic (or geometric) average* computes the aggregate confidence in each class as a weighted algebraic (or geometric) average of the individual confidences in that class (Jacobs, 1995; Tax et al., 1997). Some variants of boosting employ weighted average combination (Drucker et al., 1993).

## 3. Notation

Let  $\mathbf{A} = \langle A_1, A_2, \dots, A_L \rangle$  denote a vector of  $L$  attribute variables with domain  $\mathbf{D} = D_1 \times \dots \times D_L$ . Denote a corresponding vector of values (i.e., instantiated variables) as  $\mathbf{a} = \langle a_1, a_2, \dots, a_L \rangle \in \mathbf{D}$ . Each vector  $\mathbf{a}$  is categorized into one of  $M$  *classes*,  $C_1, C_2, \dots, C_M$ . There are  $N$  *classifiers*, or learners, which attempt to learn a functional mapping from instantiated attributes to classes. Different types of classifiers return different amounts of information—some return a single vote for one predicted class, others return a ranking of the classes, and still others return confidence scores for all classes.<sup>2</sup> Our contention is that confidence information is usually available, whether explicitly (e.g., from neural net activation values, or Bayesian net or decision tree likelihoods) or implicitly from observed performance on the training data. Thus we denote learner  $i$ ’s *classification* as an assignment  $\langle S_{i1}, \dots, S_{iM} \rangle$  of confidence scores to the classes, where  $S_{ij} \in \mathbb{R}$ . Each classifier is a function  $f_i : \mathbf{D} \rightarrow \mathbb{R}^M$ . When confidence magnitude information is truly unavailable, we adopt Lee and Srihari’s (1995) conventions for encoding classifications: A single vote for class  $C_j$  is represented as a classification vector with a 1 in the  $j$ th position and zeros elsewhere; a rank list of the classes is represented as a vector with a 1

<sup>2</sup>These three output conditions correspond to Lee and Srihari’s (1995) definitions of Type I, Type II, and Type III classifiers, respectively.

in the top class position,  $1 - 1/M$  in the second place position,  $1 - 2/M$  in the third place position, etc. Note that, technically, these two encodings introduce unfounded comparative information. For example, a vote for  $C_j$  conveys only that all other classes are less preferred than  $C_j$ , but are otherwise incomparable among themselves. Variants of the limitative theorems in this paper are also possible using more faithful representations of votes and rankings.

An *ensemble combination function*  $g$  accepts an  $N$ -tuple of classifications and returns a composite classification; that is,  $g : \mathbf{K} \rightarrow \Re^M$ , where  $\mathbf{K} \subseteq (\Re^M)^N$ . Thus, assuming  $\mathbf{K} = (\Re^M)^N$ , the aggregate classification of arbitrary classifiers  $f_1, \dots, f_N$  on an input  $\mathbf{a}$  is  $g(f_1(\mathbf{a}), \dots, f_N(\mathbf{a}))$ .

For a given input vector  $\mathbf{a} \in \mathbf{D}$ , we find it convenient to define  $\mathbf{S}$  as the  $N \times M$  matrix of all learners' confidence scores for all classes. That is,  $S_{ij}$  is learner  $i$ 's confidence that  $\mathbf{a}$  is in class  $j$ . Let  $\mathbf{r}_i$  be an  $N$ -dimensional row vector with a 1 in the  $i$ th position and zeros elsewhere; similarly, let  $\mathbf{c}_j$  be an  $M$ -dimensional column vector with a 1 in the  $j$ th position and zeros elsewhere. Then  $\mathbf{r}_i\mathbf{S}$  is the  $i$ th row of  $\mathbf{S}$ , and  $\mathbf{Sc}_j$  is the  $j$ th column of  $\mathbf{S}$ . In other words,  $\mathbf{r}_i\mathbf{S} = f_i(\mathbf{a})$  is learner  $i$ 's classification, and  $\mathbf{Sc}_j$  is the vector of all confidence scores for class  $j$ . Note that  $\mathbf{r}_i\mathbf{Sc}_j = S_{ij}$ . We denote the ensemble classification by  $\mathbf{S}_0 = \langle S_{01}, S_{02}, \dots, S_{0M} \rangle = g(\mathbf{S})$ . We write  $\mathbf{v} > \mathbf{w}$  to indicate that every component of  $\mathbf{v}$  is strictly greater than the corresponding component of  $\mathbf{w}$ .

## 4. Multiple Classes

In this section, we propose a normative basis for ensemble learning when  $M \geq 3$ . Our treatment is similar in spirit to Pennock, Horvitz and Giles's (2000) analysis of the axiomatic foundations of collaborative filtering.

### 4.1 An Impossibility Theorem

We present five properties adopted from Social Choice theory, argue their merits in the context of ensemble learning, and describe which existing algorithms exhibit which properties. Each property places a constraint on the allowable form of  $g$ .

**Property 1 (Universal domain (UNIV))**  $\mathbf{K} = (\Re^M)^N$ .

UNIV requires that  $g$  be defined for any combination of classification vectors. Since an arbitrary classifier may return an arbitrary classification, it seems only reasonable that  $g$  should return some result in all circumstances. All existing ensemble combination methods, to our knowledge, are defined for all possible classifier output patterns.

**Property 2 (Non-dictatorship (ND))** *There is no dictator  $i$  such that, for all classification matrices  $\mathbf{S}$  and all classes  $j$  and  $k$ ,  $S_{ij} > S_{ik} \Rightarrow S_{0j} > S_{0k}$ .*

In words,  $g$  is not permitted to completely ignore all but one of the classifiers, irrespective of  $\mathbf{S}$ . We consider the desirability of this axiom to be self-evident, since the whole point of ensemble learning is to improve upon the performance of the individual classifiers.

**Property 3 (Weak Pareto principle (WP))** *For all classes  $j$  and  $k$ ,  $\mathbf{Sc}_j > \mathbf{Sc}_k \Rightarrow S_{0j} > S_{0k}$ .*

WP captures the natural ideal that, if *all* classifiers are strictly more confident about one class than another, then this relationship should be reflected in the ensemble classification. Essentially all voting schemes (e.g., plurality, pairwise majority, Borda count) satisfy WP. Weighted plurality and weighted averaging methods obey WP when all weights are nonnegative (and at least one is positive). If a particular classifier's predictions are bad enough, some combination functions (e.g., weighted average with negative weights, or stacking) may establish a negative dependence between that classifier's opinion and the ensemble result, and thus violate WP. However, researchers typically strive to generate ensembles of algorithms that are as accurate as possible for a given amount of diversity (Dietterich, 1997; Dietterich, in press).

**Property 4 (Independence of irrelevant alternatives (IIA))** *Consider two classification matrices  $\mathbf{S}, \mathbf{S}'$ . If  $\mathbf{Sc}_j = \mathbf{S}'\mathbf{c}_j$  and  $\mathbf{Sc}_k = \mathbf{S}'\mathbf{c}_k$ , then  $S'_{0j} > S'_{0k} \Leftrightarrow S_{0j} > S_{0k}$ .*

Under IIA, the final relative ranking between two classes cannot depend on the confidence scores for any other classes. For example, suppose that, in classifying a fruit as either an apple, a banana, or a pear, the ensemble concludes that "apple" is most likely. Now imagine that we learn one piece of categorical knowledge (and nothing else): the fruit is *not* a pear. Every classifier diminishes its confidence in "pear", but leaves its relative confidences between "apple" and "banana" untouched. Intuitively, the ensemble should not suddenly conclude that the fruit is a banana; indeed, admitting such a reversal is contrary to most formal reasoning procedures, including Bayesian reasoning. Seemingly unfounded reversals like this are precisely what IIA guards against. Weighted averaging methods do satisfy IIA, although plurality vote, and most other voting techniques, can violate it. In Section 7, we illustrate the paradoxical results than can occur when IIA is not met.

**Property 5 (Scale invariance (SI))** *Consider two classification matrices  $\mathbf{S}, \mathbf{S}'$ . If  $\mathbf{r}_i\mathbf{S}' = \alpha_i\mathbf{r}_i\mathbf{S} + \beta_i$  for all  $i$  and for any positive constants  $\alpha_i$  and any constants  $\beta_i$ , then  $S'_{0j} > S'_{0k} \Leftrightarrow S_{0j} > S_{0k}$  for all classes  $j$  and  $k$ .*

Different classifiers (especially those based on different learning algorithms) may report confidences using different scales—one, say, ranging from 0 to 1; another from -100 to 100. Even if they share a common range, one classifier may tend to report confidence scores in the high end

of the scale, while another tends to use the low end. SI reflects the intuition that all classifiers' scores should be normalized to a common scale before combining them. One natural normalization is:

$$\mathbf{r}_i \mathbf{S}' \leftarrow \frac{\mathbf{r}_i \mathbf{S} - \min(\mathbf{r}_i \mathbf{S})}{\max(\mathbf{r}_i \mathbf{S}) - \min(\mathbf{r}_i \mathbf{S})}. \quad (1)$$

This transforms all confidence scores to the  $[0, 1]$  range, filtering out any dependence on multiplicative ( $\alpha_i$ ) or additive ( $\beta_i$ ) scale factors.<sup>3</sup> Lee and Srihari justify a similar normalization simply because “each output [classification] vector is defined over a different space” (1995, p.42). Ensemble combination schemes based on votes or rankings are by definition invariant to scale; weighted averaging methods, on the other hand, are not.

Different researchers favor differing subsets of these five properties, at least implicitly via their choice of combination methods. Roberts (1980) proves that *no* combination algorithm whatsoever can “have it all”.

**Proposition 1 (Impossibility)** *If  $M > 2$ , no function  $g$  simultaneously satisfies UNIV, ND, WP, IIA, and SI.*

**Proof:** Follows from Sen's (1986) or Roberts's (1980, Theorem 3) extensions of Arrow's (1963) original theorem. ■

## 4.2 Weighted Average Combination

We might weaken SI, allowing the final classification to depend on the magnitudes of confidence differences, but not on additive scale shifts.

**Property 6 (Translation invariance (TI))** *Consider two classification matrices  $\mathbf{S}, \mathbf{S}'$ . If  $\mathbf{r}_i \mathbf{S}' = \alpha \mathbf{r}_i \mathbf{S} + \beta_i$  for all  $i$  and for any (single) positive constant  $\alpha$  and any constants  $\beta_i$ , then  $S'_{0j} > S'_{0k} \Leftrightarrow S_{0j} > S_{0k}$  for all classes  $j$  and  $k$ .*

TI can be enforced by an additive normalization, or aligning all classifiers' scores with a common reference point (e.g.,  $\mathbf{r}_i \mathbf{S}' \leftarrow \mathbf{r}_i \mathbf{S} - \min(\mathbf{r}_i \mathbf{S})$ ).

This weakening is sufficient to allow for a non-dictatorial combination function  $g$ . Moreover, the only such  $g$  computes the ensemble confidence in each class as a weighted average of the component learners' confidences in that class.

**Proposition 2 (Weighted Average)** *If  $M > 2$ , then the only function  $g$  satisfying UNIV, WP, IIA, and TI is such that  $\mathbf{w} \mathbf{S} c_j > \mathbf{w} \mathbf{S} c_k \Rightarrow S_{0j} > S_{0k}$ , where  $\mathbf{w} = \langle w_1, w_2, \dots, w_N \rangle$  is a row vector of  $N$  nonnegative weights, at least one of which is positive. If  $g$  is also continuous, then  $\mathbf{w} \mathbf{S} c_j > \mathbf{w} \mathbf{S} c_k \Leftrightarrow S_{0j} > S_{0k}$ .*

<sup>3</sup>If  $\max(\mathbf{r}_i \mathbf{S}) = \min(\mathbf{r}_i \mathbf{S})$  then set  $\mathbf{r}_i \mathbf{S}'$  to 0.

**Proof:** Follows from Roberts's (1980) Theorem 2. ■

Certainly there may exist classification domains where some of these properties do not seem appropriate or justified. However, we believe that, because the properties are very natural, understanding the limitations that they place on the space of ensemble learning algorithms helps to clarify what potential algorithms can and cannot do.

## 5. Binary Classification

Now consider the subset of learning problems where  $M = |C| = 2$ . In this case, the impossibility outlined in Proposition 1 disappears; the five properties UNIV, WP, IIA, SI, and ND are in fact perfectly compatible. For example, all five are satisfied by the standard *majority vote*:

$$\|S_{01} - S_{02}\| \equiv \left\| \sum_{i=1}^N \|S_{i1} - S_{i2}\| \right\| \quad (2)$$

where

$$\|x\| = \begin{cases} 1 & : \text{if } x > 0 \\ 0 & : \text{if } x = 0 \\ -1 & : \text{if } x < 0 \end{cases}.$$

Note that the properties are necessary but not sufficient for characterizing majority vote. Proposition 3 below provides one sufficient characterization.

### 5.1 Majority Vote

The use of majority vote for ensemble learning is typically motivated by its simplicity, its observed effectiveness, and its perceived fairness when the constituent algorithms are essentially “created equal” (Dietterich, 1997). For example, the component algorithms employed for bagging, ECOC, and randomization are generally *a priori* indistinguishable, and (2) is typically used to combine classifications in these cases.

May (1952) provides an *axiomatic* justification for majority vote. His treatment is directly applicable when the constituent algorithms return only votes (equivalent to rankings since  $M = 2$ ), rather than arbitrary confidence scores. We now generalize his axioms and his characterization theorem to apply to confidence scores.

**Property 7 (Neutrality (NTRL))**

$$\begin{aligned} \text{If } g(\langle S_{11}, S_{12} \rangle, \dots, \langle S_{N1}, S_{N2} \rangle) &= \langle S_{01}, S_{02} \rangle \\ \text{then } g(\langle S_{12}, S_{11} \rangle, \dots, \langle S_{N2}, S_{N1} \rangle) &= \langle S_{02}, S_{01} \rangle. \end{aligned}$$

Under NTRL, the effect of every algorithm reversing its vote is simply to reverse the aggregate vote. NTRL establishes a symmetry between the two class names,  $C_1$  and  $C_2$ , ruling out any *a priori* bias for one class name over the

other. Indeed, the subscripts 1 and 2 are assigned to the two classes arbitrarily; NTRL simply ensures that the final result does not depend on how the two classes are indexed. NTRL is a strictly stronger constraint than IIA.

#### Property 8 (Symmetry (SYM))

$$\begin{aligned} g(\langle S_{11}, S_{12} \rangle, \dots, \langle S_{N1}, S_{N2} \rangle) \\ = g(\langle S_{i_11}, S_{i_12} \rangle, \dots, \langle S_{i_N1}, S_{i_N2} \rangle) \end{aligned}$$

where  $\{i_1, i_2, \dots, i_N\}$  is any permutation of  $\{1, 2, \dots, N\}$ .

SYM is stronger than ND and is sometimes referred to as *anonymity*. Whereas NTRL implies an invariance under class name reversal, SYM enforces an invariance under any permutation of algorithm names, or subscripts. It simply insists that our numbering scheme has no effect on the output of the combination rule. Note that SYM *does not*, by itself, rule out a posterior bias based on the classifiers' reported confidence scores.

**Property 9 (Positive Responsiveness (POSR))** Consider two classification matrices  $\mathbf{S}, \mathbf{S}'$ . If  $\|S_{01} - S_{02}\| \in \{0, 1\}$ , and  $\mathbf{r}_i \mathbf{S}' = \mathbf{r}_i \mathbf{S}$  for all  $i \neq h$ , and  $\mathbf{r}_h \mathbf{S}'$  is such that either

1.  $S'_{h1} > S_{h1}$  and  $S'_{h2} = S_{h2}$ , or
2.  $S'_{h1} = S_{h1}$  and  $S'_{h2} < S_{h2}$ ,

then  $\|S'_{01} - S'_{02}\| = 1$ .

If the current aggregate vote is tied ( $\|S_{01} - S_{02}\| = 0$ ), then, under POSR, any change by any algorithm  $i$  in a positive direction for  $C_1$  (i.e.,  $S_{h1}$  increases or  $S_{h2}$  decreases) breaks this deadlock, yielding  $S_{01} > S_{02}$ . Moreover, any change of one of the constituent votes that strictly favors  $C_1$  cannot swing the ensemble vote in the opposite direction, from  $C_1$  to undecided or to  $C_2$ . Combined with NTRL, POSR is a stronger version of WP, but is still quite reasonable. Note that, because there are only two classes, if any learner's votes are observed to be negatively correlated with the correct classification (and, for example, a weighted average method assigns a negative weight), then its votes can simply be reversed, rendering POSR (and a nonnegative weight) appropriate again.

**Proposition 3 (Majority Vote)** An aggregation function  $g$  is the majority vote (2) if and only if it satisfies UNIV, SI, NTRL, SYM, and POSR.

**Proof:** Choose scaling parameters as in Equation 1:  $\alpha_i = (|S_{i1} - S_{i2}|)^{-1}$  (or if  $S_{i1} = S_{i2}$ , set  $\alpha_i = 1$ ) and  $\beta_i = -\alpha_i \min(S_{i1}, S_{i2})$ . Let  $\mathbf{r}_i \mathbf{S}' = \alpha_i \mathbf{r}_i \mathbf{S} + \beta_i$  for all  $i$ . Then

$$\langle S'_{i1}, S'_{i2} \rangle = \begin{cases} \langle 1, 0 \rangle & : \text{if } S_{i1} > S_{i2} \\ \langle 0, 0 \rangle & : \text{if } S_{i1} = S_{i2} \\ \langle 0, 1 \rangle & : \text{if } S_{i1} < S_{i2} \end{cases}.$$

That is, with only two classes, and two degrees of freedom in choosing the scaling constants, SI effectively restricts the domain  $\mathbf{K}$  of  $g$  to votes. May (1952) proves that NTRL, SYM, and POSR are necessary and sufficient conditions for majority vote when inputs are votes. We refer the reader to May's article for the remainder of the proof. ■

Notice that, when the component algorithms return only votes, and no other information is available, SI is a vacuous requirement; in this setting, Proposition 3 becomes a very compelling normative argument for the use of majority vote for classifier combination.

#### 5.2 Weighted Majority Vote

When the component algorithms do return meaningful confidence scores, SI may seem overly severe, as it essentially strips away magnitude information. Confidence scores may reflect many sources of information—for example, the activation levels of a neural network's output nodes, the posterior probabilities of a Bayesian network's output variables, or an algorithm's observed performance on the training data (as is used in Boosting). Regardless of its origin we interpret  $S_{i1} > S_{i2}$  as a prediction in favor of class one,  $S_{i2} > S_{i1}$  as a prediction in favor of class two, and the magnitude of the difference in confidence scores  $|S_{i2} - S_{i1}|$  as the *weight* of algorithm  $i$ 's conviction.

Then we define the *weighted majority vote* as

$$\begin{aligned} \|S_{01} - S_{02}\| &\equiv \left\| \sum_{i=1}^N |S_{i1} - S_{i2}| \cdot \|S_{i1} - S_{i2}\| \right\| \\ &= \left\| \sum_{i=1}^N S_{i1} - S_{i2} \right\|. \end{aligned} \quad (3)$$

#### Property 10 (Separable Symmetry (SSYM))

$$\begin{aligned} g(\langle S_{11}, S_{12} \rangle, \dots, \langle S_{N1}, S_{N2} \rangle) \\ = g(\langle S_{i_11}, S_{j_12} \rangle, \dots, \langle S_{i_N1}, S_{j_N2} \rangle) \end{aligned}$$

where  $\{i_1, i_2, \dots, i_N\}$  and  $\{j_1, j_2, \dots, j_N\}$  are any two permutations of  $\{1, 2, \dots, N\}$ .

SSYM is a stronger constraint than SYM. Under SSYM, the ensemble classification depends on the set of confidence scores for class one and the set of confidence scores for class two, but not on the identity of the algorithms that return those scores.

**Proposition 4 (Weighted Majority Vote)** The only aggregation function  $g$  that satisfies UNIV, TI, NTRL, SSYM, and POSR is the weighted majority vote (3).

**Proof:** Under UNIV and NTRL,  $\mathbf{S} = \mathbf{0}$  implies that  $S_{01} = S_{02}$ . Thus, under POSR, if  $S_{N1} > S_{N2}$  and

$S_{i1} = S_{i2} = 0$  for all  $i \neq N$ , then  $S_{01} > S_{02}$ . Similarly, because of NTRL, if  $S_{N2} > S_{N1}$  and  $S_{i1} = S_{i2} = 0$  for all  $i \neq N$ , then  $S_{02} > S_{01}$ . Given an arbitrary classification matrix  $\mathbf{S}$ , we can make the following invariance transformations. We invoke TI and SSYM alternately and repeatedly as follows:

$$\begin{aligned} g(\langle\langle S_{11}, S_{12} \rangle, \langle S_{21}, S_{22} \rangle, \langle S_{31}, S_{32} \rangle, \dots \rangle) &= \\ g(\langle\langle S_{11} - S_{12}, 0 \rangle, \langle 0, S_{22} - S_{21} \rangle, \langle S_{31}, S_{32} \rangle, \dots \rangle) &= \\ g(\langle\langle 0, 0 \rangle, \langle S_{11} - S_{12}, S_{22} - S_{21} \rangle, \langle S_{31}, S_{32} \rangle, \dots \rangle) &= \\ g(\langle\langle 0, 0 \rangle, \langle S_{11} + S_{21} - S_{12} - S_{22}, 0 \rangle, \\ &\quad \langle 0, S_{32} - S_{31} \rangle, \dots \rangle) = \\ \dots &= \\ g\left(\left\langle\langle 0, 0 \rangle, \langle 0, 0 \rangle, \langle 0, 0 \rangle, \dots, \left\langle\sum_{i=1}^N S_{i1} - S_{i2}, 0\right\rangle\right\rangle\right) \end{aligned}$$

Thus if  $\sum_i S_{i1} - S_{i2}$  is greater than (less than, equal to) zero, then  $S_{01} - S_{02}$  is greater than (less than, equal to) zero, precisely the weighted majority vote (3). ■

## 6. Independence Preservation

Consider the learners' predictions when asked to evaluate an example  $\mathbf{a}^*$  with some missing values. Without loss of generality, let  $A_1, A_2, \dots, A_m$  be the attribute variables with missing values, and let  $A_{m+1}, \dots, A_L$  be the variables with known values. Let  $\mathbf{a}_{m+}^* = \langle a_{m+1}^*, \dots, a_L^* \rangle$  denote the vector of known values. If we define a prior joint probability distribution  $\Pr(\mathbf{a})$  over all possible combinations of attribute values, then we can compute each learner's induced posterior distribution over classifications given the known values  $\mathbf{a}_{m+}^*$ :

$$\Pr_i(\mathbf{r}_i \mathbf{S} | \mathbf{a}_{m+}^*) = \sum_{\substack{\mathbf{x} \in \{D_1 \times \dots \times D_m\}: \\ f_i(\mathbf{x}, \mathbf{a}_{m+}^*) = \mathbf{r}_i \mathbf{S}}} \Pr(\mathbf{x} | \mathbf{a}_{m+}^*).$$

Similarly, we can compute the ensemble's posterior distribution over classifications:

$$\Pr_0(\mathbf{S}_0 | \mathbf{a}_{m+}^*) = \sum_{\substack{\mathbf{x} \in \{D_1 \times \dots \times D_m\}: \\ g(f_1(\mathbf{x}, \mathbf{a}_{m+}^*), \dots, f_N(\mathbf{x}, \mathbf{a}_{m+}^*)) = \mathbf{S}_0}} \Pr(\mathbf{x} | \mathbf{a}_{m+}^*).$$

Now we can ascertain whether some attributes are statistically independent of the classification. Again without loss of generality, select attribute  $A_{m+1}$  for this purpose. What if *every* constituent algorithm agrees that  $A_{m+1}$  is independent of the classification, given the remaining known values  $a_{m+2}^*, \dots, a_L^*$ ? It seems natural and desirable that such a unanimous judgment of "irrelevance" should be preserved in the ensemble distribution. The following property formally captures this ideal:

$A_1$	$A_2$	$A_3$	$\mathbf{r}_1 \mathbf{S}$	$\mathbf{r}_2 \mathbf{S}$	$\mathbf{r}_3 \mathbf{S}$	$\mathbf{S}_0$
0	0	0	$\langle 1, 0 \rangle$	$\langle 0, 1 \rangle$	$\langle 1, 0 \rangle$	$\langle 1, 0 \rangle$
0	1	0	$\langle 0, 1 \rangle$	$\langle 1, 0 \rangle$	$\langle 0, 1 \rangle$	$\langle 0, 1 \rangle$
1	0	0	$\langle 1, 0 \rangle$	$\langle 0, 1 \rangle$	$\langle 0, 1 \rangle$	$\langle 0, 1 \rangle$
1	1	0	$\langle 1, 0 \rangle$	$\langle 1, 0 \rangle$	$\langle 1, 0 \rangle$	$\langle 1, 0 \rangle$
$\Pr_i(\langle 1, 0 \rangle   A_3 = 0)$			0.75	0.5	0.5	0.5

$A_1$	$A_2$	$A_3$	$\mathbf{r}_1 \mathbf{S}$	$\mathbf{r}_2 \mathbf{S}$	$\mathbf{r}_3 \mathbf{S}$	$\mathbf{S}_0$
0	0	1	$\langle 0, 1 \rangle$	$\langle 0, 1 \rangle$	$\langle 1, 0 \rangle$	$\langle 0, 1 \rangle$
0	1	1	$\langle 1, 0 \rangle$	$\langle 0, 1 \rangle$	$\langle 1, 0 \rangle$	$\langle 1, 0 \rangle$
1	0	1	$\langle 1, 0 \rangle$	$\langle 1, 0 \rangle$	$\langle 0, 1 \rangle$	$\langle 1, 0 \rangle$
1	1	1	$\langle 1, 0 \rangle$	$\langle 1, 0 \rangle$	$\langle 0, 1 \rangle$	$\langle 1, 0 \rangle$
$\Pr_i(\langle 1, 0 \rangle   A_3 = 1)$			0.75	0.5	0.5	0.75

Table 1. Example where plurality vote violates IPP.

### Property 11 (Independence Preservation Property (IPP))

If  $\Pr_i(\mathbf{r}_i \mathbf{S} | \mathbf{a}_{m+}^*) = \Pr_i(\mathbf{r}_i \mathbf{S} | a_{m+2}^*, \dots, a_L^*)$  for all  $i$   
then  $\Pr_0(\mathbf{S}_0 | \mathbf{a}_{m+}^*) = \Pr_0(\mathbf{S}_0 | a_{m+2}^*, \dots, a_L^*)$ .

Table 1 presents a constructive proof that plurality vote fails to satisfy IPP. Three attributes each have domain  $D_i = \{0, 1\}$ , and the prior distribution over attribute values  $\Pr(\mathbf{a}) = 1/8$  is uniform. Variables  $A_1$  and  $A_2$  have missing values (i.e.,  $m = 2$ ). Each of three constituent algorithms agree that the classification is independent of  $A_3$ . But combination by plurality vote destroys this independence: According to the ensemble, the classification *does* in fact depend on the value of  $A_3$ . Similar examples demonstrate that algebraic and geometric averages also violate IPP. It remains an open question whether any reasonable ensemble combination function can satisfy IPP. Results from Statistics concerning generalized variants of IPP are mostly negative: No acceptable aggregation function has been found that preserves independence (Genest & Zidek, 1986), and several impossibility theorems severely restrict the space of potential candidates (Genest & Wagner, 1987; Pennock & Wellman, 1999).

## 7. Experimental Observations

We have shown, in theory, that the class of potential ensemble algorithms is severely limited if we want a small number of intuitive properties satisfied. One might argue that situations where these properties come into conflict may never arise in practice if we use popular aggregation methods. The purpose of this section is to show by example that, in fact, such conflicts do occur in practice. Specifically, we will give examples from a stock market prediction domain where IIA breaks down if we base our aggregation on voting.

rank order	#	rank order	#
UP > SAME > DOWN	6	DOWN > SAME > UP	5
UP > DOWN > SAME	1	SAME > UP > DOWN	5
DOWN > UP > SAME	3	SAME > DOWN > UP	1

Table 2. Six learned vote patterns, and the number of neural networks that learned each. An instance of the Borda paradox.

To this end, we report results of empirical tests of an ensemble learner trained on stock market data. We retrieved daily closing prices of the Dow between 1/20/97 and 1/18/00 from MSN Investor.<sup>4</sup> From this, we generated an approximately zero-mean and unit-variance time series of the form  $\{d_t = 85(\ln p_t - \ln p_{t-1})\}$ , where  $p_t$  is the Dow's price on day  $t$ . The attributes are  $\mathbf{A} = \langle d_{t-5}, d_{t-4}, \dots, d_{t-1} \rangle$ . The classes are discrete intervals of  $d_t$  such that  $C_1 = \text{UP} \equiv (d_t > 0.35)$ ,  $C_2 = \text{DOWN} \equiv (d_t < -0.35)$ , and  $C_3 = \text{SAME} \equiv (-0.35 \leq d_t \leq 0.35)$ . The intervals are such that each class frequency is roughly 1/3. The component learning algorithms are backpropagation neural networks built using Flake's (1999) NODELIB code library; each consists of an input layer of five nodes, a hidden layer of from one to seven nodes, and an output layer of three nodes. Diversity is due only to differences in the number of hidden nodes and to randomization in the training algorithm. The time series  $d_t$  was divided into a training set of 562 days and a test set of 187 days.

Table 2 shows the learned class rankings for twenty one networks (three each with 1, 2, ..., 7 hidden nodes) on test day 7/14/99. If we use standard plurality vote to combine predictions, then DOWN wins with 8 votes, UP places in second with 7 votes, and SAME comes in last with 6 votes. By this measure we should short the Dow. But are we sure? Since SAME is presumably the least likely outcome, let's focus on the relative likelihoods between only DOWN and UP.<sup>5</sup> If we ignore SAME and recompute the vote, we find that UP actually beats DOWN by 12:9! This is a vivid demonstration that plurality vote violates IIA; the preference between UP and DOWN depends on SAME. So should we invest in the Dow? Well, the other two pairwise majority votes reveal that SAME beats UP by 11:10 and SAME beats DOWN by 12:9. Then according to the pairwise majority, SAME wins against both other classes, UP comes in second, and DOWN is last, completely reversing the original order predicted by the three-way plurality vote. This is an illustration of the so-called *Borda voting paradox*, named after the eighteenth century scientist who discovered it.

Table 3 demonstrates another classic voting paradox, due to Condorcet, one of Borda's peers. The table lists the ac-

$i$	$S_{i1}$	$S_{i2}$	$S_{i3}$	rank order
1	-0.33	-0.41	-0.25	SAME > UP > DOWN
2	-0.45	-0.25	-0.27	DOWN > SAME > UP
3	-0.31	-0.35	-0.37	UP > DOWN > SAME

Table 3. Confidence scores and corresponding vote patterns for three neural networks. An instance of the Condorcet paradox.

tivation values (confidence scores) of three networks (with one, two, and three hidden nodes) on test day 4/23/99. Plurality vote is tied, since each algorithm ranks a different class highest. What about pairwise majority vote? In this case, SAME beats UP by 2:1, and UP beats DOWN by 2:1. So is SAME our predicted outcome? Not necessarily—DOWN beats SAME, also by 2:1. We see that pairwise majority vote can return cyclical predictions, a violation of our generic definition of a classification  $\mathfrak{R}^M$ , which assumes that aggregation returns a transitive ordering of classes.

These two “paradoxes” illustrate the undesirable consequences of violating some of the basic properties of  $g$  defined earlier. The examples also constitute an existence proof that some of the same counterintuitive outcomes that have perplexed social scientists for centuries can and do occur in the context of ensemble learning.

## 8. Conclusion

We identified several properties of combination functions that Social Choice theorists and statisticians have found compelling, and argued their applicability in the context of ensemble learning. We cataloged common ensemble methods according to the properties they do and do not satisfy, and showed that no combination function can possess them all. We provided axiomatic justifications for weighted average combination, majority vote, and weighted majority vote. We described how common aggregation methods fail to respect unanimous judgments of independence. Finally, we exemplified the fundamental and unavoidable tradeoffs among the various properties using an ensemble learner trained on stock market data.

Drucker, et al. (1993) present empirical evidence that weighted average outperforms plurality vote in some circumstances. Future work will examine whether the axiomatic framework developed in this paper can aid in deriving theoretical bounds on the performance of weighted average and other combination rules. We also plan to explore normative justifications for *individual* classifiers, and investigate whether, in some cases, a complex individual classifier might reasonably be interpreted as an ensemble of simpler constituent classifiers.

<sup>4</sup><http://moneycentral.msn.com/investor>

<sup>5</sup>Or we may have received outside information that discounts the likelihood of SAME.

## Acknowledgements

Thanks to Gary Flake, Michael Wellman, and the anonymous reviewers. Pedrito Maynard-Reid II was partially supported by a National Physical Science Consortium Fellowship.

## References

- Arrow, K. J. (1963). *Social choice and individual values*. New Haven, CT: Yale University Press. 2nd edition.
- Bauer, E., & Kohavi, R. (1999). An empirical comparison of voting classification algorithms: Bagging, boosting, and variants. *Machine Learning*, 36, 105–142.
- Breiman, L. (1996). Bagging predictors. *Machine Learning*, 24, 123–140.
- Dietterich, T. G. (1997). Machine learning research: Four current directions. *AI Magazine*, 18, 97–136.
- Dietterich, T. G. (in press). An experimental comparison of three methods for construction ensembles of decision trees: Bagging, boosting, and randomization. *Machine Learning*.
- Dietterich, T. G., & Bakiri, G. (1995). Solving multiclass learning problems via error-correcting output codes. *Journal of Artificial Intelligence Research*, 2, 263–286.
- Drucker, H., Schapire, R., & Simard, P. (1993). Boosting performance in neural networks. *International Journal of Pattern Recognition and Artificial Intelligence*, 7, 704–719.
- Flake, G. W. (1999). *An industrial strength numerical modeling research and development tool* (Technical Report 99-085). NEC Research Institute.
- Freund, Y., & Schapire, R. E. (1996). Experiments with a new boosting algorithm. *Proceedings of the Thirteenth International Conference on Machine Learning* (pp. 148–156).
- Freund, Y., & Schapire, R. E. (1999). A short introduction to boosting. *Journal of the Japanese Society for Artificial Intelligence*, 14, 771–780.
- Genest, C., & Wagner, C. G. (1987). Further evidence against independence preservation in expert judgement synthesis. *Aequationes Mathematicae*, 32, 74–86.
- Genest, C., & Zidek, J. V. (1986). Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*, 1, 114–148.
- Hansen, L., & Salamon, P. (1990). Neural network ensembles. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12, 993–1001.
- Jacobs, R. A. (1995). Methods for combining experts' probability assessments. *Neural Computation*, 7, 867–888.
- Jordan, M. I., & Jacobs, R. A. (1994). Hierarchical mixtures of experts and the EM algorithm. *Neural Computation*, 6, 181–214.
- Lee, D., & Srihari, S. N. (1995). A theory of classifier combination: The neural network approach. *Proceedings of the Third International Conference on Document Analysis and Recognition* (pp. 42–45).
- Madhvanath, S., & Govindaraju, V. (1995). Serial classifier combination for handwritten word recognition. *Proceedings of the Third International Conference on Document Analysis and Recognition* (pp. 911–914).
- May, K. O. (1952). A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica*, 20, 680–684.
- Pennock, D. M., Horvitz, E., & Giles, C. L. (2000). Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering. *Proceedings of the Seventeenth National Conference on Artificial Intelligence*. In press.
- Pennock, D. M., & Wellman, M. P. (1999). Graphical representations of consensus belief. *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence* (pp. 531–540).
- Roberts, K. W. S. (1980). Interpersonal comparability and social choice theory. *Review of Economic Studies*, 47, 421–439.
- Schapire, R. E. (1990). The strength of weak learnability. *Machine Learning*, 5, 197–227.
- Sen, A. (1986). Social choice theory. In K. J. Arrow and M. D. Intriligator (Eds.), *Handbook of mathematical economics*, vol. iii, 1073–1181. Amsterdam: Elsevier.
- Tax, D. M. J., Duin, R. P. W., & van Breukelen, M. (1997). Comparison between product and mean classifier combination rules. *Proceedings of the Workshop on Statistical Pattern Recognition*.
- Tumer, K., & Ghosh, J. (1996). Error correlation and error reduction in ensemble classifiers. *Connection Science*, 8, 385–404.
- Wolpert, D. (1992). Stacked generalization. *Neural Networks*, 5, 241–260.

---

## Aggregating Learned Probabilistic Beliefs

---

**Pedrito Maynard-Reid II**  
 Computer Science Department  
 Stanford University  
 Stanford, CA 94305-9010  
 pedmayn@cs.stanford.edu

**Urszula Chajewska**  
 Computer Science Department  
 Stanford University  
 Stanford, CA 94305-9010  
 urszula@cs.stanford.edu

### Abstract

We consider the task of aggregating beliefs of several experts. We assume that these beliefs are represented as probability distributions. We argue that the evaluation of any aggregation technique depends on the semantic context of this task. We propose a framework, in which we assume that nature generates samples from a ‘true’ distribution and different experts form their beliefs based on the subsets of the data they have a chance to observe. Naturally, the optimal aggregate distribution would be the one learned from the combined sample sets. Such a formulation leads to a natural way to measure the accuracy of the aggregation mechanism.

We show that the well-known aggregation operator LinOP is ideally suited for that task. We propose a LinOP-based learning algorithm, inspired by the techniques developed for Bayesian learning, which aggregates the experts’ distributions represented as Bayesian networks. We show experimentally that this algorithm performs well in practice.

## 1 Introduction

Belief aggregation of subjective probability distributions has been a subject of great interest in statistics (see [GZ86, CW99]) and, more recently, artificial intelligence (e.g., [PW99]) and machine learning (ensemble learning in particular [PMGH00]), especially since probabilistic distributions are increasingly being used in medicine and other fields to encode knowledge of experts. Unfortunately, many of the aggregation proposals have lacked sufficient semantical underpinnings, typically evaluating a mechanism by how well it satisfies properties justified by little more than intuition. However, as has been noted in other fields such as belief revision (cf. [FH96]), the appropriateness of properties depends on the particular context.

We take a more semantic approach to aggregation: we first describe the realistic framework in which the experts or *sources* learn their probability distributions from data using standard probabilistic learning techniques. We assume

a Decision Maker (DM) — the traditional name for the aggregator — wants to aggregate a set of these learned distributions. This framework suggests a natural optimal aggregation mechanism: construct the distribution that would be learned had all the sources’ data sets been available to the DM. Since the original data sets are generally not available, the aggregation mechanism should come as close as possible to reconstructing the data sets and learning from the combined set.

For intuition, consider the task of creating an expert system for some specialized medical field. We would like to take advantage of the expertise of several doctors working in this field. Each of these doctors sharpened his knowledge by following many patients. The doctors can no longer recall the specifics of each case, but they have formed over the years fairly accurate models of the domain that can be represented as sets of conditional probabilities. (In fact, many expert systems have been created over the years by eliciting such conditional probabilities from experts [HHN92].) Of course, if there was a doctor who had seen all of the patients the others doctors saw, the ideal expert system would result from eliciting her model. However, there isn’t one such expert. Therefore, our system would benefit from incorporating the knowledge of as many experts as we can find. The system would also account for the differing levels of experience of different doctors – some of them may have practiced for much longer than others.

One of the best-known aggregation operators is the Linear Opinion Pool (LinOP) which aggregates a set of distributions by taking their weighted sum. It has been shown in the statistics community that, under some intuitive assumptions, learning the joint distribution from the combined data set is equivalent to using LinOP over the individual joint distributions learned from the individual data sets. However, whereas the weights in typical uses of LinOP are often criticized for being ad-hoc, our framework prescribes semantically-justified weights: the estimated percentages of the data each source saw. Intuitively, a high weight means we believe a source has seen a relatively

large amount of data and is, hence, likely to be reliable.

However, joint distributions are hardly the preferred representation for probabilistic beliefs in real-world domains. BNs (aka belief networks, etc.) [Pea88] have gained much popularity as structured representations of probability distributions. They allow such distributions to be represented much more compactly, therefore often avoiding exponential blowup in both memory size and inference complexity.

Thus, we assume the sources beliefs are BNs learned from data. According to our semantics, the aggregate BN should be one the DM would learn from the combined sets of data. We describe a LinOP-based BN aggregation algorithm, inspired by the algorithm designed to learn BNs from data. The algorithm uses sources' distributions instead of samples to search over possible BN structures and parameter settings. It takes advantage of the marginalization property of LinOP to make computation more efficient. We explore the algorithm's behavior by running experiments on the well-known, real-life Alarm network [BSCC89] and on the smaller artificial Asia network [LS88].

## 2 Formal Preliminaries

We restrict our attention to domains with discrete variables. We consider how to compute the aggregate distribution, and how the accuracy of our computation depends on how much we know about the sources.

Formally, we consider the following setting: There are  $L$  sources and  $N$  discrete random variables, where each variable  $X$  has domain  $\text{dom}(X)$ . We follow the convention of using capital letters to denote variables and lowercase letters to denote their values. Symbols in bold denote sets.  $\mathcal{W}$  is the set of possible worlds defined by value assignments to variables. The true distribution or model of the world is  $\pi$ . Each source  $i$  has a data set  $\mathbf{D}_i$  sampled from (unknown to us)  $\pi$ . We will assume that each  $\mathbf{D}_i$  is finite of size  $M_i$ . The corresponding empirical (i.e., frequency) distribution is  $\hat{p}_i$ . Each source  $i$  learns a distribution  $p_i$  over  $\mathcal{W}$ . This is  $i$ 's model of the world. The combined set of samples is  $\mathbf{D} = \cup_i \mathbf{D}_i$  of size  $M$ . The corresponding empirical distribution is  $\hat{p}$ . The DM constructs an aggregate distribution  $p$ . The optimal aggregate distribution  $p^*$  is posited to be the distribution the DM would learn from  $\mathbf{D}$ .

Since it is unrealistic to expect the DM to have access to the sources' sample sets, we consider how to use information about the sources' learned distributions to at least approximate  $p^*$ . Specifically, we consider the situation where the DM knows the sources' distributions and has a good estimate of the percentage  $\alpha_i = M_i/M$  of the combined set of samples each source  $i$  has observed as well as what learning method it used.

We make a number of assumptions. First, we assume that the samples are not noisy or otherwise corrupted, and they

are complete (no missing values).

Second, we assume that the individual sample sets are disjoint (so  $M = \sum_i M_i$ ). This implies that the concatenation of the  $\mathbf{D}_i$  equals  $\mathbf{D}$ , so we don't have to concern ourselves with repeats when aggregating. This assumption is not always appropriate. It is invalidated when multiple sources observe the same event. However, there are interesting domains where this property holds. For example, in our motivating medical domain, doctors are likely to have seen disjoint sets of patients.

Third, we assume that the sources believe their samples to be IID — independent and identically distributed. The machine learning algorithms used in practice commonly rely on this assumption.

Finally, we assume that the samples in the combined set  $\mathbf{D}$  are sampled from  $\pi$  and IID. This assumption may appear overly restrictive at first glance. For one, it may seem to preclude the common situation where sources receive samples from different subpopulations. For example, if doctors are in different parts of the world, the characteristics of the patients they see will likely be different.

In fact, we can accommodate this situation within our framework by assuming  $\pi$  is a distribution over the domain variables and a source variable  $S$  which takes the different sources as values;  $S = i$  means source  $i$  observed the instantiated domain variables. This generalized distribution is sampled IID. Each  $\mathbf{D}_i$  consists of the subset of samples where  $S = i$ . It is not necessary to keep around the  $S$  values; computing the  $p_i$  and  $p^*$  without  $S$  will give the same results as learning distributions over the complete samples and marginalizing out  $S$ . Thus, although samples will be IID, different subpopulation distributions will be possible, captured by different conditional probability distributions of the domain variables given distinct values of  $S$ .<sup>1</sup>

## 3 Aggregating Learned Joint Distributions

We first consider the case where sources have learned joint distributions, and the aggregate is also a joint.

### 3.1 Learning joint distributions: review

Given samples of a variable  $X$ , the goal of a learner is to estimate the probability of future occurrences of each value of  $X$ . In our setting, the domain of  $X$  is  $\mathcal{W}$  and the parameters that need to be learned are the  $|\mathcal{W}|$  probabilities. The distribution over  $X$  is parameterized by  $\Theta$ . Two standard approaches are *Maximum Likelihood Estimation (MLE)* and *Maximum A Posteriori estimation (MAP)*.

---

<sup>1</sup>Two implications of this formulation are that the assumption that the  $\mathbf{D}_i$  are disjoint is implicit and  $\alpha_i$  will approach  $\pi(S = i)$  as  $M$  approaches  $\infty$  for all  $i$ .

An MLE learner chooses the member of a specified family of distributions that maximizes the likelihood of the data:

**Definition 1** If  $X$  is a random variable,  $\text{dom}(X) = \{x_1, \dots, x_k\}$ , and  $\Theta = \langle \Theta_1, \dots, \Theta_k \rangle$  where  $\Theta_i = P(x_i | \Theta)$ , then the MLE distribution over  $X$  given data set  $\mathbf{D}$  is

$$\text{MLE}(X, \mathbf{D}) = \arg \max_{\theta} P(\mathbf{D} | \theta)$$

It is easy to show that the MLE distribution is the empirical distribution if samples are IID.

MAP learning, on the other hand, follows the Bayesian approach to learning which directs us to put a prior distribution over the value of any parameter we wish to estimate. We treat these parameters as random variables and define a probability distribution over them. More formally, we now have a joint probability space that includes both the data and the parameters.

**Definition 2** If  $X$  is a random variable,  $\text{dom}(X) = \{x_1, \dots, x_k\}$ , and  $\Theta = \langle \Theta_1, \dots, \Theta_k \rangle$  where  $\Theta_i = P(x_i | \Theta)$ , then the MAP distribution over  $X$  given data set  $\mathbf{D}$  and prior  $P(\Theta)$  is the distribution

$$\text{MAP}(X, P(\Theta), \mathbf{D}) \equiv P(X | \mathbf{D}) = \int P(X | \Theta)P(\Theta | \mathbf{D}) d\Theta$$

The appropriate conjugate prior for variables with multinomial distributions is Dirichlet.  $\text{Dir}(\Theta | \gamma_1, \dots, \gamma_k)$ , where each  $\gamma_i$  is a hyperparameter such that  $\gamma_i > 0$ .

We will assume that Dirichlet distributions are assessed using the *method of equivalent samples*: given a prior distribution  $\rho$  over  $X$  and an estimated sample size  $\xi$ ,  $\gamma_i$  is simply  $\rho(x_i)\xi$ . We use these to parameterize MAP:

**Definition 3** If  $X$  is a random variable,  $\text{dom}(X) = \{x_1, \dots, x_k\}$ ,  $\Theta = \langle \Theta_1, \dots, \Theta_k \rangle$  where  $\Theta_i = p(x_i | \Theta)$ ,  $\rho$  is a probability distribution over  $X$ , and  $\xi > 0$ , then  $\text{MAP}(X, \langle \rho, \xi \rangle, \mathbf{D})$  denotes the distribution  $\text{MAP}(X, p_0, \mathbf{D})$  where  $p_0 = \text{Dir}(\Theta | \rho(x_1)\xi, \dots, \rho(x_k)\xi)$ .

We will omit the  $X$  argument from the MLE and MAP notation since it is understood.

### 3.2 LinOP: review

Let us turn to the problem of aggregation. We will show that joint aggregation essentially reduces to LinOP. LinOP was proposed by Stone in [Sto61], but is generally attributed to Laplace. It aggregates a set of joint distributions by taking a weighted sum of them:

**Definition 4** Given probability distributions  $p_1, \dots, p_L$  and non-negative parameters  $\beta_1, \dots, \beta_L$  such that  $\sum_i \beta_i = 1$ , the LinOP operator is defined such that, for

any  $w \in \mathcal{W}$ ,

$$\text{LinOP}(\beta_1, p_1, \dots, \beta_L, p_L)(w) = \sum_i \beta_i p_i(w).$$

LinOP is popular in practice because of its simplicity. As described in [GZ86], it also has a number of attractive properties such as *unanimity* (if all the  $p_i = p'$ , then LinOP returns  $p'$ ), *non-dictatorship* (no one input is always followed), and the *marginalization property* (aggregation and marginalization are commutative operators). However, LinOP has often been dismissed in the aggregation communities as a normative aggregation mechanism, primarily because it fails to satisfy a number of other properties deemed to be necessary of any reasonable aggregator, e.g., the *external Bayesianity* property (aggregation and conditioning should commute) and the preservation of shared independences. Furthermore, typical approaches to choosing the weights are often criticized as being ad-hoc.

However, this dismissal may have been overly hasty. LinOP proves to be the operator we are looking for in our framework: using it is equivalent to having the DM learn from the combined data set under intuitive assumptions.

### 3.3 MLE aggregation

Suppose the sources and the DM are MLE learners. As has been known in statistics for some time, the DM need only compute the LinOP of the sources' distributions.

**Proposition 1 ([Win68, Mor83])** If  $p_i = \text{MLE}(\mathbf{D}_i)$  for each  $i \in \{1, \dots, L\}$  and  $p^* = \text{MLE}(\mathbf{D})$ , then  $p^* = \text{LinOP}(\alpha_1, p_1, \dots, \alpha_L, p_L)$ .

Although straight-forward, this proposition is illuminating. For one, the weight corresponding to each source has a very clear meaning: it is the percentage of total data seen by that source. The DM only needs to provide accurate estimates of these percentages. A high weight indicates that the DM believes a source has seen a relatively large amount of data and is, hence, likely to be very reliable. Thus, we address a common criticism of LinOP, that the weights are often chosen in an ad-hoc fashion. Also, if  $M$  is known, the DM can compute the number of samples in  $\mathbf{D}$  that were  $w$ :  $\text{MLinOP}(\alpha_1, p_1, \dots, \alpha_L, p_L)$ . Thus, LinOP can be viewed as essentially storing the *sufficient statistics* for the DM learning problem.

It is now easy to see why a property such as preservation of independence will not always hold given our learning-based semantics. In our framework, sources do not have strong beliefs about independences; any believed independence depends on how well it fits the source's data. The independence preservation property does not take into account the possibility that, because of limited data, sources may all have learned independences which are not justified if all the data was taken into account.

Consider, for example, the following distribution over two variables  $A$  and  $B$ :  $\pi(ab) = 1/4$ ,  $\pi(a\bar{b}) = 1/6$ ,  $\pi(\bar{a}b) = 1/3$ , and  $\pi(\bar{a}\bar{b}) = 1/4$ . Obviously,  $A$  and  $B$  are not independent. Suppose two sources have each received a set of six samples from this distribution:  $\mathbf{D}_1$  consists of one each of  $ab$  and  $a\bar{b}$ , two each of  $\bar{a}b$  and  $\bar{a}\bar{b}$ ;  $\mathbf{D}_2$  consists of one each of  $\bar{a}b$  and  $\bar{a}\bar{b}$ , two each of  $ab$  and  $\bar{a}b$ . Further suppose each used MLE to learn a distribution over  $A$  and  $B$ .  $A$  and  $B$  are independent in each of these distributions. The LinOP distribution, on the other hand, effectively takes into account the evidence seen by both sources and actually computes  $\pi$  where the variables are not independent.

### 3.4 MAP aggregation

MLE learners are known to have problems with overfitting and low-probability events for which data never materialized. MAP learning often does a better job of dealing with these problems, especially when data is sparse.

Consequently, suppose the sources and the DM are MAP learners with Dirichlet priors. The optimal aggregate distribution is a variation on LinOP:<sup>2</sup>

**Proposition 2** Suppose, for each  $i \in \{1, \dots, L\}$ ,  $p_i = \text{MAP}(\langle p_i, \xi_i \rangle, \mathbf{D}_i)$  and  $p^* = \text{MAP}(\langle \rho, \xi \rangle, \mathbf{D}_i)$ . Then,

$$p^*(w) = \frac{1}{M + \xi} (M \text{LinOP}(\alpha_1, p_1, \dots, \alpha_L, p_L) + \rho(w)\xi) + \sum_i \frac{\xi_i}{M + \xi} (p_i(w) - \rho_i(w)). \quad (1)$$

The first term in Equation 1 is the DM's MAP estimation, the second term accounts for the sources' priors by subtracting out their effect.

**Corollary 2.1** Suppose, for each  $i \in \{1, \dots, L\}$ ,  $p_i = \text{MAP}(\langle p_i, \xi_i \rangle, \mathbf{D}_i)$  and  $p^* = \text{MAP}(\langle \rho, \xi \rangle, \mathbf{D}_i)$ . Then,

$$\lim_{\xi/M \rightarrow 0} p^* = \text{LinOP}(\alpha_1, p_1, \dots, \alpha_L, p_L).$$

Thus, as  $M$  becomes large, the LinOP distribution approaches  $p^*$ . This is not surprising since it is well-known that MLE learning and MAP learning with Dirichlet priors are asymptotically equivalent. The implication is that if  $M$  is large, not only do we not need to know  $M$  to aggregate, we do not need to know what priors the sources used either. And if we approximate the aggregate distribution by the LinOP distribution, this approximation will improve the more samples seen by the sources.

## 4 Aggregating Learned Bayesian Networks

Bayesian networks (BNs) are structured representations of probability distributions. A BN  $b$  consists of a directed

<sup>2</sup>We omit proofs for lack of space.

acyclic graph (DAG)  $g$  whose nodes are the  $N$  random variables. The parents of a node  $X$  are denoted by  $\text{Pa}(X)$ ;  $\text{pa}(X)$  denotes a particular assignment to  $\text{Pa}(X)$ . The structure of the network encodes marginal and conditional independencies present in the distribution. Associated with each node is the conditional probability distribution (CPD) for  $X$  given  $\text{Pa}(X)$ .

We consider the case where sources' beliefs are represented as BNs learned from data. We briefly review the techniques used for learning BNs from data. For a more detailed presentation, see [Hec96].

### 4.1 Learning Bayesian networks: review

If the structure of the network is known, the task reduces to statistical parameter estimation by MLE or MAP. In the case of complete data, the likelihood function for the entire BN conveniently decomposes according to the structure of the network, so we can maximize the likelihood of each parameter independently.

If the structure of the network is not known, we have to apply Bayesian model selection. More precisely, we define a discrete variable  $G$  whose states  $g$  correspond to possible models, i.e., possible network structures; we encode our uncertainty about  $G$  with the probability distribution  $P(g)$ . For each model  $g$ , we define a continuous vector-valued variable  $\Theta_g$ , whose instantiations  $\theta_g$  correspond to the possible parameters of the model. We encode our uncertainty about  $\Theta_g$  with a probability distribution  $P(\theta_g | g)$ .

We score the candidate models by evaluating the *marginal likelihood* of the data set  $\mathbf{D}$  given the model  $g$ , that is, the *Bayesian score*  $P(\mathbf{D} | g) = \int P(\mathbf{D} | \theta_g, g)P(\theta_g | g)P(g)d\theta_g$ .

In practice, we often use some approximation to the Bayesian score. The most commonly used is the MDL score, which converges to the Bayesian score as the data set becomes large. The MDL score is defined as

$$\begin{aligned} \text{score}_{\text{MDL}}(b' : \mathbf{D}) &= \\ & M \sum_{i=1}^N \sum_{\text{pa}(X_i)} \sum_{x_i} \hat{p}(x_i, \text{pa}(x_i)) \log \hat{p}(x_i | \text{pa}(x_i)) \\ & - \frac{\log M}{2} \text{Dim}[g'] - \text{DL}(g') \end{aligned}$$

where  $\text{Dim}[g']$  is the number of independent parameters in the graph and  $\text{DL}(g')$  is the description length of  $g'$ . Finding the network structure with the highest score has been shown to be NP-hard in general. Thus, we have to resort to heuristic search. Since the search can easily get stuck in a local maximum, we often add random restarts to the process. The BN learning algorithm is presented in Figure 1.

Why are we interested in learning BNs rather than joint

```

1. pick a random DAG g
2. parameterize g to form b
3. score b
4. loop
5.   for each DAG g' differing from g by
        adding, removing, or reversing an edge
6.     parameterize g' to form b'
7.     score b'
8.   pick the b' with the highest score and replace
        g with g' and b with b' if score(b') > score(b)
9. until no further change g
10. return b.

```

Figure 1: Bayesian network learning algorithm.

distributions? Besides some obvious reasons concerning compact representation and efficient inference, a distribution learned by the BN algorithm may be closer to the original distribution used to generate the data in the first place.

First, note that the networks which can be parameterized to represent exactly the MLE- or MAP-learned joint distributions are, in general, fully connected. Intuitively, a distribution learned from finite sample data will always be a little noisy, so true independences will almost always look like slight dependences mathematically. As a result, the BNs we are interested in (either for the sources or for the DM) will not be exact representations of the independencies present in the MLE- or MAP-learned distributions, but, rather, will account for this overfitting.

BN learning ‘stretches’ the distribution that best fits the data to match candidate network structures. For every structure, we look for the best (producing the highest score) parameterization *of that structure*. The score balances the fit to the data with model complexity.

## 4.2 LinOP-based Aggregation Algorithm

Now suppose each source has learned a BN  $b_i$  with DAG  $g_i$  from  $\mathbf{D}_i$  using the MDL score and the DM is given these BNs as well as the  $\alpha_i$ . According to our semantics, the aggregate BN should be as close as possible to the one the DM would learn from  $\mathbf{D}$ .

We cannot apply the BN learning algorithm directly, since we don’t have the data used by sources to learn their models. A simple solution would be to generate samples from each source model and train the DM on the combined set. That algorithm, although appealingly simple, raises some new questions. It is not clear how many samples we should generate from each source. One possibility would be to use the same number as the (estimated) number of samples that each source used to learn its model. However, if that number is small, the samples will not represent the generating distribution adequately, introducing additional noise to the process. If we generate more samples than each source saw (increasing it proportionally to preserve the  $\alpha_i$  settings), we give too much weight to the MLE component of the score, thus possibly choosing a suboptimal network. In fact, our

experiments described in Section 5 show that this algorithm does very badly in practice.

Instead, we can adapt the BN learning algorithm to use sources’ distributions instead of samples.

The main difference is in the way we compute the MLE/MAP parameters for each structure we consider and the way we compute the score (lines 2, 3, 6 and 7 in Figure 1). Our algorithm relies on the observation that it is not necessary to have the actual data to learn a BN; it is sufficient to have their empirical distribution. As we have demonstrated in Section 3, we can come up with said distribution by applying the LinOP operator to distributions learned by our sources.

We can take advantage of the marginalization property of LinOP to make computation more efficient. As is noted in [PW99], we can parameterize the network in top-down fashion by first computing the distribution over the roots, then joints over the second layer variables together with their parents, etc. The conditional probabilities can be computed by dividing the appropriate marginals (using Bayes Law). In many cases, that would require only local computations in sources’ BNs.

The MDL score also requires knowing only the empirical distribution for  $\mathbf{D}$  and  $M$ . Again, since the empirical distribution is the LinOP distribution if the weights are chosen correctly and the sources used MLE or MAP (assuming sufficient data) learning, it is possible to score the candidate networks without having the actual data. Furthermore, the marginals used in the MLE score are family marginals. If the previous parameterization step is done by computing marginals, then these will have already been computed.

Although the MDL score requires knowledge of  $M$ , this dependence may not be strong, especially for large  $M$  in which case the second term is dominated by the likelihood term and  $M$  becomes a factor common to all networks and can be ignored. Otherwise, a rough approximation of  $M$  should suffice.

As in traditional BN learning, caching can make the parameterization and scoring of ‘neighboring’ networks more efficient. Since we are making only local changes to the structure, only a few parameters will need updating. If an arc is added or removed, we only need to recompute new parameters for the child node, and if an arc is switched, we only need to recompute parameters for the two nodes involved. Also, since these LinOP marginals don’t change, caching computed values may help to further speed up future computations.

## 5 Experiments

We implemented the BN aggregation algorithm in Matlab using Kevin Murphy's Bayes Net Toolbox<sup>3</sup> and explored its behavior by running experiments on the well-known, real-life Alarm network [BSCC89], a 37-node network used as part of a system for monitoring intensive care patients, and on the smaller 8-node artificial Asia network [LS88].

In our experiments, we learned two source BNs from data sampled from the original BN, then aggregated the results using our algorithm (AGGR). We had both the sources and the DM use MAP to parameterize their networks. In computing LinOP, we used the  $\alpha_i$  as weights. We compared our proposal's accuracy against learning from the combined data sets (OPT) by plotting the Kullback-Leibler (KL) divergence [Kul59]<sup>4</sup> of each distribution from the true distribution for different values of  $M = |\mathbf{D}|$ .

### 5.1 Sensitivity to $M$

We considered the situation where the DM knows the priors used by the sources and adjusts for the unduly large number of imaginary samples. All sources and DMs used the Dirichlet prior defined by the uniform distribution and an estimated sample size of 1. We varied the total number of samples  $M$  between 200 and 20000, having sources see the same number of samples in some cases and different numbers in others. We conducted multiple runs for each setting and averaged them. Figure 2(a) plots the averages for the Alarm network when sources have equal  $\alpha_i$ . Due to software limitations, we had to start each structure search with the fully disconnected graph and used no random restarts for this larger network. As can be seen, in spite of the limited search, our algorithm does fairly well as far as coming close to the optimal and improving on the sources. Not surprisingly, the KL divergence drops as the total number of samples increases. Furthermore, the experiments on sources with different  $\alpha_i$  showed no dependence of the performance of the algorithm on the relative difference in  $\alpha_i$ .

We ran similar experiments on Asia. Here, we varied the number of samples between 200 and 3000, with five runs per setting. For each run, we used five random restarts. Figure 2(b) plots the average for each setting. The plot shows that when we are able to explore the search space sufficiently in the learning and aggregation algorithms, our algorithm consistently improves on the sources and closely approximates to the optimal.

<sup>3</sup>Available at <http://www.cs.berkeley.edu/~murphyk/bnt.html>.

<sup>4</sup>The KL divergence of distribution  $q$  from  $p$  is defined as  $\sum_{w \in \mathcal{W}} p(w) \log \frac{p(w)}{q(w)}$ .

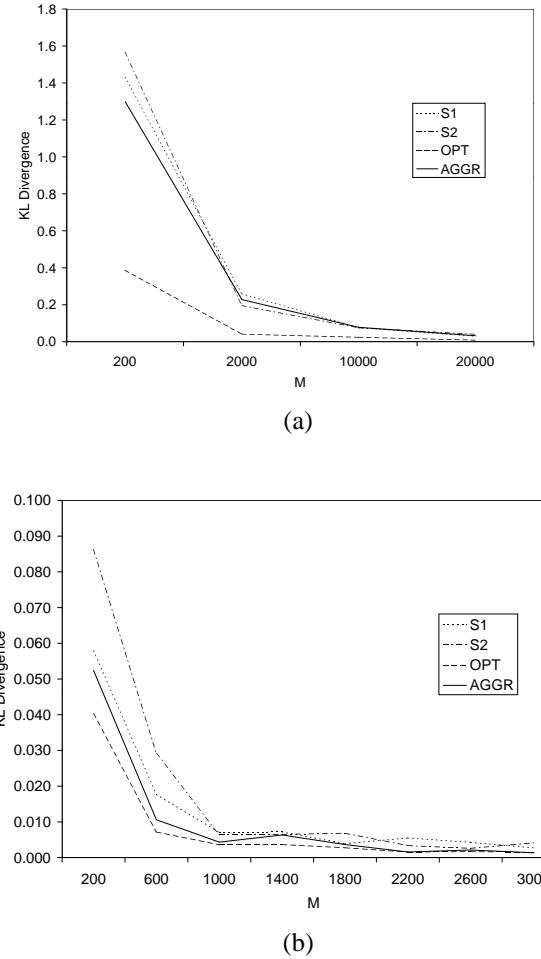


Figure 2: Sensitivity to  $M$  (a) Alarm network results. (b) Asia network results.

### 5.2 Sensitivity to the DM's estimation of $M$

We hypothesized earlier that the actual value of the DM's estimate of  $M$  does not matter all that much. To demonstrate this, we ran experiments on the Asia network similar to those above, but leaving  $M$  fixed and varying the DM's estimate 1 order of magnitude above and below  $M$ . Figures 3(a) summarizes the results for  $M = 100$ .

Any approximation above 0.25 orders of magnitude below  $M$  provides improvement over the sources. Estimates below this made the complexity penalty sufficiently strong to select DAGs with fewer arcs than the original and underfit the data. On the other hand, although overestimating  $M$  did not increase the KL distance from the original, there is a danger of extreme overestimates causing overfitting. However, we did not find any increase in the complexity of the aggregate networks for the 1 order of magnitude range we considered; they remained at 8–9 arcs on average.

Figure 3(b) summarizing the results for  $M = 10000$  shows

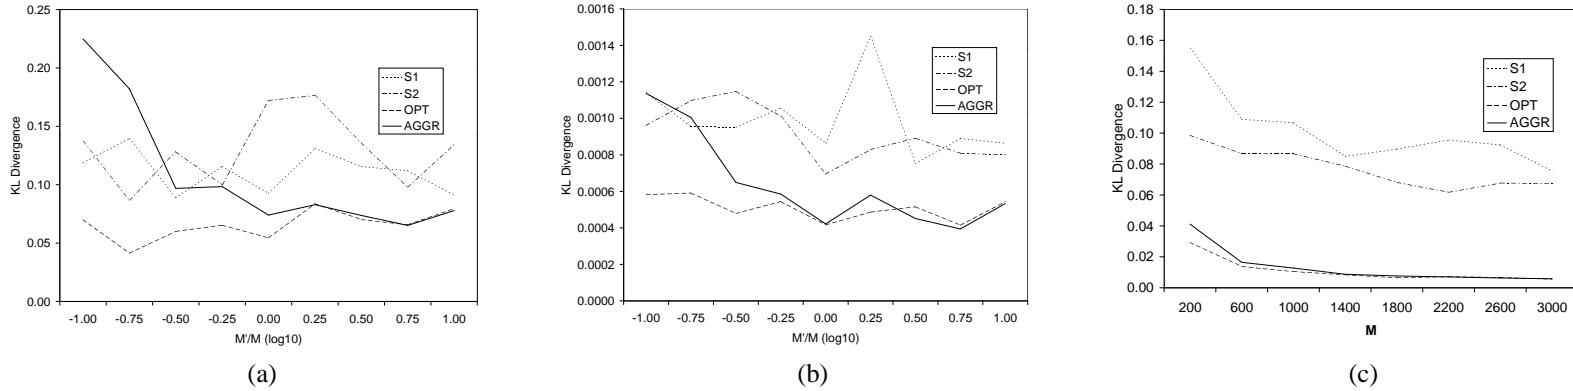


Figure 3: Asia network results (a) varying DM’s estimate of  $M$  ( $M = 100$ ). (b) varying DM’s estimate of  $M$  ( $M = 10k$ ). (c) with different subpopulations.

that, as predicted, the range of “slack” increases with  $M$ ; the more samples seen by the sources, the less important the accuracy of the DM’s estimate.

### 5.3 Subpopulations

Our algorithm performs well when combining source distributions learned based on samples from different subpopulations. To show this, we modified the Asia network to accomodate two sources, a doctor practicing in San Francisco and one practicing in Cincinnati. The probability distributions of the two root nodes in the Asia network, representing whether a patient smokes and whether she has visited Asia would be significantly different for the two doctors. A patient from San Francisco is less likely to be a smoker, and one from Cincinnati is less likely to have visited Asia. Thus, we added a source variable as described in Section 2, gave the sources equal priors of seeing patients, made the source variable a parent of the two root variables, and gave them appropriate CPDs. We drew  $M$  samples from this extended network and had each source learn from the appropriate subset, then used AGGR to combine the results using the correct  $\alpha_i$  and  $M$ . Figure 3(c) plots the KL divergence of each distribution from the original distribution with the source variable marginalized out. Because the sources are learning the distributions for different subpopulations, what they learn is relatively far from the overall distribution. The DM takes advantage of the information from both sources and learns a BN that approximates the original much more closely than either source.

### 5.4 Comparison to sampling algorithm

In each of the above experiments, we also compared the performance of our algorithm to the alternative intuitive algorithm SAMP we described in Section 4.2 in which we sample  $\alpha_i M$  samples from each source  $i$ ’s BN and learn a BN from the combined data. SAMP did very badly in

general, consistently worse than not only AGGR, but worse than the sources as well, often by an order of magnitude.

## 6 Related Work

A wealth of work exists in statistics on aggregating probability distributions. Good surveys of the field include [GZ86, CW99]. Many of the earlier, axiomatic approaches suffered from a lack of semantical grounding. For this reason, the community moved towards modeling approaches instead. The most studied approach has been the *supra-Bayesian* one, introduced in [Win68] and formally established in [Mor74, Mor77]. Here, the DM has a prior not only over the variables in the domain, but over the possible beliefs of the sources as well. She aggregates by using Bayesian conditioning to incorporate the information she receives from the sources. In fact, Proposition 1 derives from this body of work. However, almost all of this work has been restricted to aggregating beliefs represented as point probabilities or odds, or joint distributions.

There has been some recent interest, particularly in AI, in the problem of aggregating structured distributions including [MA92, MA93, PW99]. But, like the early axiomatic approaches in statistics, much of this work focuses on attempting to satisfy abstract properties such as preserving shared independences, and often runs into impossibility results as a consequence.

In some sense, what we are doing could also be viewed as ensemble learning for BNs. Ensemble learning involves combining the results of different weak learners to improve classification accuracy. Because of its simplicity, LinOP is often used without justification to do the actual combination. Our results justify this use when the weak learners use MLE, MAP, or BN learning.

Another new area in AI that bears similarities to our work is that of on-line or incremental learning of BNs (e.g.,

[Bun91, LB94, FG97]). There, we are given a continuous stream of samples and we want to maintain a BN learned from all the data we have seen so far. Because the stream is very long, it is generally not possible to maintain the full set of sufficient statistics. Approaches range from approximating the sufficient statistics to restricting the network that can be learned. We essentially do the former by assuming that the sufficient statistics for the data seen by each source is encoded in its network. Cross-fertilization between the two fields may prove profitable.

## 7 Conclusion

We have presented a new approach to belief aggregation. We believe that we cannot formulate that problem precisely or measure success of different techniques without answering questions about the way in which sources' beliefs were formulated. We argued that a framework in which the sources are assumed to have learned their distributions from data is both intuitively plausible and leads to a very natural formulation of the optimal DM distribution — one which would be learned from the combined data sets — and a natural success measure — a distance from the generating, ‘true’ distribution.

Based on the observation that LinOP is the appropriate operator for this framework if sources and DM are MLE learners, we presented a LinOP-based algorithm to aggregate beliefs represented by Bayesian networks. Our preliminary results show that this algorithm performs very well.

One direction of future work will involve finding ways to relax the various assumptions. For example, we would like to extend the framework to allow for continuous variables and to allow for dependence between sources' sample sets.

In our framework, the DM completely ignores sources' priors. This may be appropriate if the priors are known to be unreliable or uninformative. However, the priors used in real applications are often informative in and of themselves. Thus, a second direction will involve finding valid ways of taking advantage of sources' priors to improve the quality of the aggregation. For example, if sources use Dirichlet priors and the DM trusts their estimated sample sizes, she may chose to incorporate them into her estimate of  $M$ .

## Acknowledgements

Pedrito Maynard-Reid II was partially supported by a National Physical Science Consortium Fellowship. Urszula Chajewska was supported by the Air Force contract F30602-00-2-0598 under DARPA's TASK program.

## References

- [BSCC89] I. Beinlich, G. Suermondt, R. Chavez, and G. Cooper. The ALARM monitoring system. In

- Proc. European Conf. on AI and Medicine*, 1989.
- [Bun91] W. Buntine. Theory refinement on bayesian networks. In *Proc. UAI'91*, pages 52–60, 1991.
- [CW99] R. T. Clemen and R. L. Winkler. Combining probability distributions from experts in risk analysis. *Risk Analysis*, 19(2):187–203, 1999.
- [FG97] N. Friedman and M. Goldsmith. Sequential update of bayesian network structure. In *Proc. UAI'97*, pages 165–174, 1997.
- [FH96] N. Friedman and J. Y. Halpern. Belief revision: A critique. In *Proc. KR'96*, pages 421–431, 1996.
- [GZ86] C. Genest and J. V. Zidek. Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*, 1(1):114–148, 1986.
- [Hec96] D. Heckerman. A tutorial on learning bayesian networks. Technical Report MSR-TR-95-06, Microsoft Research, 1996.
- [HHN92] D. Heckerman, E. Horvitz, and B. Nathwani. Toward normative expert systems: Part I. The Pathfinder project. *Methods of Information in Medicine*, 31:90–105, 1992.
- [Kul59] S. Kullback. *Information Theory and Statistics*. Wiley, 1959.
- [LB94] W. Lam and F. Bacchus. Learning bayesian belief networks: An approach based on the mdl principle. *Computational Intelligence*, 10:269–293, 1994.
- [LS88] S. L. Lauritzen and D. J. Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems. In *J. Royal Statistical Society, Series B (Methodological)*, volume 50(2), pages 157–224, 1988.
- [MA92] I. Matzkevich and B. Abramson. The topological fusion of Bayes nets. In *Proc. UAI'92*, pages 191–198, 1992.
- [MA93] I. Matzkevich and B. Abramson. Some complexity considerations in the combination of belief networks. In *Proc. UAI'93*, pages 152–158, 1993.
- [Mor74] P. A. Morris. Decision analysis expert use. *Management Science*, 20:1233–1241, 1974.
- [Mor77] P. A. Morris. Combining expert judgements: A bayesian approach. *Management Science*, 23:679–693, 1977.
- [Mor83] P. A. Morris. An axiomatic approach to expert resolution. *Management Science*, 29(1):24–32, 1983.
- [Pea88] J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, 1988.
- [PMGH00] D. M. Pennock, P. Maynard-Reid II, C. L. Giles, and E. Horvitz. A normative examination of ensemble learning algorithms. In *Proc. ICML'00*, pages 735–742, 2000.
- [PW99] D. M. Pennock and M. P. Wellman. Graphical representations of consensus belief. In *Proc. UAI'99*, pages 531–540, 1999.
- [Sto61] M. Stone. The opinion pool. *Annals of Mathematical Statistics*, 32(4):1339–1342, 1961.
- [Win68] Robert L. Winkler. The consensus of subjective probability distributions. *Management Science*, 15(2):B61–B75, October 1968.